

# Evaluating an Energy-efficient Radio Architecture for Opportunistic Communication

Sylvia T. Kouyoumdjieva, Ólafur Helgason, Emre A. Yavuz, and Gunnar Karlsson

Laboratory for Communication Networks

KTH, Electrical Engineering

Stockholm, Sweden

{stkou, olafurr, emreya, gk}@ee.kth.se

**Abstract**—A major challenge for wireless networks is to minimize the energy consumption in the mobile devices. This work evaluates potential performance gains of a dual-radio architecture in which a low power radio is used to wake up the primary high power radio. Our targeted domain is *opportunistic communication* directly between devices. We show that it can significantly reduce the energy consumption at a price of slightly decreased goodput. We also examine the effect of the MAC protocol on the performance of the dual-radio system and we point out that in dense scenarios MAC layer protocols, such as of 802.15.4, do not perform well. We observe that information on device density of an environment can be used to address this shortcoming.

## I. INTRODUCTION

Mobile wireless networks without an infrastructure (such as ad hoc networks and opportunistic networks) have been proposed as a feasible way for sharing contents among mobile devices [1], [2]. Although these networks possess certain interesting properties and are considered as a viable mechanism for offloading some types of data traffic in infrastructure systems [3], one of the key challenges which still needs to be addressed is the energy consumption. Since content distribution in mobile wireless networks is mainly targeted at mobile users with battery powered devices, it is of great importance that the protocols and mechanisms used are energy efficient. Although recent smartphones are powerful devices with advanced networking and multimedia capabilities, battery capacity is still a scarce resource. In particular, the wireless LAN interface is responsible for a large fraction of the energy consumption [4].

We have confirmed this in our implementation of an opportunistic content distribution system for the Android platform, using 802.11 radio interface in ad-hoc mode [5]. Figure 1 shows how the 802.11 radio interface (in ad-hoc mode) affects the battery life on a HTC Hero smartphone. When the interface is turned on, the battery life is reduced to only 25% of what it is with the interface turned off. This is despite the fact that no data has been transmitted or received via the interface during the test. Once the interface is turned on, it consumes relatively high power regardless of being in a *transmit* or *receive* state since listening in *idle* state consumes almost as much energy. This suggests that reducing or eliminating the *idle* energy cost of the 802.11 interface may be a promising strategy to reduce the overall energy consumption and prolong battery life. As a




Fig. 1. The effect of 802.11 interface energy consumption on battery life.

minimum requirement, the battery of a smartphone running an opportunistic networking service should last at least a full day so that it can be recharged during night. Minimizing the energy consumption of the networking subsystem is therefore one of the key challenges for enabling opportunistic communication as a viable communication mode.

In this work we study and evaluate the potential performance gains that can be achieved by equipping mobile handhelds with a dual-radio system. A primary high-power, high-bitrate radio is used for data transfer and a secondary low-power, low-bitrate radio is used as a control channel for performing neighbor discovery. The high power data radio can thus be suspended and only woken up when a neighbor is discovered by the low-power control-channel radio. This way, the high-power radio is only used for transmission/reception of data and its idle-mode cost is significantly reduced.

It has been shown that a dual radio approach is feasible to implement and experimentally verified that it can significantly reduce energy consumption [6], [7]. These small-scale evaluations (2-3 nodes) focused on infrastructure systems where the low-power radio is used to wake up the high-power radio when it is in reach of an access point or when another node wants to communicate with the local node via the access point. In contrast, our work focuses on evaluating the efficiency of a dual radio approach in an *opportunistic setting* where no fixed infrastructure can be assumed and nodes are *not synchronized*. Our study thus incorporates (1) development of an extensive framework for simulating dual-radio architecture

in opportunistic context, (2) adoption of realistic pedestrian mobility models, and (3) presentation of a novel approach for maximizing the probability of successful data transfer via filtering of contact opportunities.

The rest of this paper is structured in the following way. Section II further motivates our work and contrasts with other related work. In section III we describe our evaluation scenario and simulation model. Section IV presents our results and in section V we conclude.

## II. BACKGROUND AND RELATED WORK

Mobile handhelds today have commonly two available radios that can operate in ad-hoc mode: Bluetooth and 802.11. Bluetooth requires less power at the expense of a shorter communication range (approx. 10 m) and lower bitrate (around 2.1 Mb/s for Bluetooth 2.0). Due to the slow scanning process for the frequency hopping spread spectrum scheme, the neighbor discovery of Bluetooth constitutes a major drawback for use in opportunistic networking. Even though the emerging Bluetooth 4.0 with its low energy feature promises shorter and faster neighbor discovery, as well as longer communication range, the lack of wide deployment prevents us from evaluating its actual performance. In comparison, 802.11 provides a higher bitrate (a raw rate of 54 Mb/s in 802.11g), longer communication range (approx. 100 m) and neighbors can quickly be discovered by using broadcast beacons. The downside is that it requires more power. However, 802.11 has a better energy-per-bit ratio than Bluetooth, making 802.11 more suitable for transmitting bulk data. In summary, 802.11 is a better candidate for opportunistic networking but its energy consumption is a major drawback.

In [8], Feeney and Nilsson measure the energy consumption of an 802.11 interface in ad-hoc mode and provide a detailed energy profile. They show that the interface consumes significant energy in idle mode and that the ad-hoc mode is not as efficient as the infrastructure mode. This is mainly because the infrastructure mode can rely on the access point to synchronize and buffer data for nodes that are in the power saving mode. In [9], Feeney and Wilkomm describe their implementation of a framework that can be used to simulate the energy consumption of a mobile node in the OMNeT++ simulator. Our evaluation on node energy consumption is based on this framework.

The *Wake on Wireless* system [6] was the first to propose a dual-radio system in the context of mobile handheld devices. Their prototype implementation used a simple low-power radio as a control channel for discovery and to wake up the 802.11 radio when a neighboring node has data to send. They show that their system can significantly increase the battery lifetime of such devices. In [7], Pering et al. further show, by prototype measurements, that significant energy can be saved by using a low-power Chipcon CC100 radio or Bluetooth for discovering 802.11 access points. Motivated by this, CoolSpots [10] proposed a dual-radio Bluetooth/802.11 architecture to switch between radio interfaces with respect to traffic intensity. Both the *Wake on Wireless* and *CoolSpots* design focuses on, and

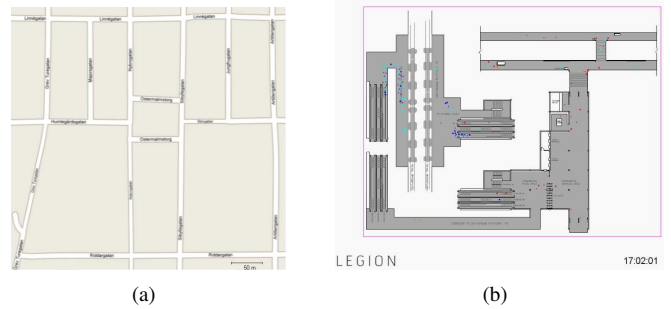


Fig. 2. The simulation scenarios: a part of downtown Stockholm (left) and a two-level subway station (right).

assume the availability of, infrastructure networks with access points. These systems are therefore not suitable for an ad-hoc scenario.

In [11], the authors consider a dual-radio system for delay-tolerant networking. They rely on GPS for clock synchronization and examine how to dimension the duty cycles of the radios for a given traffic load and contact statistics with minimal energy consumption as an objective. In contrast our work does not consider DTN routing; we assume neither GPS availability, nor reliable mobility prediction but rather focus on how incorporating application semantics in the discovery on the low-power control channel can reduce energy consumption.

## III. EVALUATION SCENARIO

We consider the power savings achievable with a dual-radio subsystem in the context of opportunistic content distribution. The users are humans in urban areas, equipped with smartphones or other handheld devices. The opportunistic content distribution system provides a publish/subscribe service, where users can subscribe to content *feeds* (i.e. topics) of their interest. Each feed contains one or more *entries*, which are the actual items of interest, such as a photoblog entry or media file. While mobile devices are in direct communication range they exchange contents matching local subscriptions. We believe that the main utility of such a system in an urban setting is when users are on the move. Hence the mobility models used in our evaluation reflect this.

For our evaluation we need to capture the system in a realistic manner on at least three levels. The *mobility scenarios* define the potential contact opportunities that arise and their frequency and duration. The *application & usage model* defines the available contents in the system, their popularity, publication rate etc. Finally, the *radio & energy model* implements the radio technologies that comprise the dual-radio subsystem (at the physical and MAC layers) and realistically models their energy consumption in various radio states.

### A. Mobility scenarios

For capturing human mobility, we use Legion Studio [12]; a commercial simulator initially developed for designing and dimensioning large-scale spaces via simulation of pedestrian

behaviors. Its multi-agent pedestrian model is based on advanced analytical and empirical models which have been calibrated by measurement studies. Legion Studio allows the use of open systems, where entities can enter and leave the system according to a predefined pattern. Each simulation run conducted in Legion Studio results in a mobility trace file, containing a snapshot of the positions of all nodes in the system every 0.6 s.

Our evaluation considers two scenarios with different characteristics: an outdoor urban scenario, modeling the Östermalm area of central Stockholm, and an indoor scenario, recreating a two-level subway station (Figure 2).

The urban outdoor scenario consists of a grid of interconnected streets. The observed area is connected to the outside world with 12 passages. The active area of the outdoor scenario is 5872 m<sup>2</sup>. The scenario can be characterized as a high mobility scenario, since nodes move constantly throughout their lifetime in the observed areas.

The indoor scenario defines a subway station with train platforms connected via escalators to the entry-level. Nodes can arrive on foot from any of the five entry points of the subway station, or when a train arrives at the platform. The train arrivals contribute to the burstiness of the node arrivals and departures. Nodes congregate while waiting for a train to arrive at one of the platforms, or while taking a break in the store or the coffee shop at the entry level. The active area of the scenario is 1921 m<sup>2</sup>.

We have chosen the input parameters of the Östermalm and the Subway scenario such that they result in approximately the same mean node density of 0.09 nodes/m<sup>2</sup> [13]. We do differentiate however the terms *node density* and *neighbor density*. While the node density gives us an idea about the average number of nodes in a space, the neighbor density represents the average number of neighbors a node has in each moment of time.

### B. Application & usage model

The application we evaluate is an opportunistic publish/subscribe content distribution system [5]. We have implemented our system for the Android platform but for the evaluation in this work we use an implementation of the same system for the OMNeT++ simulator [14].

Every node periodically broadcasts beacons to inform other nodes within communication range about its presence. When a device encounters another node, it first checks in a local cache whether it has met the node before or if so, whether the node has obtained new contents since their last contact. If this is the case, the device initiates a request-response communication and on each step of it tries to match the remote feeds/entries with its local subscriptions until it finally downloads the actual content items of interest.

When using a dual-radio approach, the beacon messages and the messages for discovering available feeds and entries are exchanged on the low power radio. Since the beacons do not contain much information, they fit into a single 802.15.4

broadcast frame. Only the actual data download is performed over the high-power radio, which is otherwise suspended.

### C. Radio & energy model

The purpose of the *radio & energy model* is to realistically model the dual-radio subsystem and its energy consumption. For this, we have adapted and extended the MiXiM framework [15] for the OMNeT++ discrete event simulator [16]. MiXiM provides detailed models of radio wave propagation, interference estimation, and wireless MAC protocols. Moreover, it incorporates the Energy Framework [9] for modeling power consumption in wireless networks. MiXiM currently only supports a single radio per device but we have extended it with an implementation of a dual-radio interface, which captures the energy consumed by different radios including their sub-states and the energy cost and delay when waking up or shutting down the high-power radio [17].

By design, MiXiM does not support inter-radio interference. Nowadays, though, most radio technologies implemented in mobile devices do operate at 2.4 GHz in the ISM band, and may in practice interfere. The radios used in a dual radio system will make no exception from this. However such inter-radio interference is currently not captured by our dual radio model. We do consider this as a possible extension of our work for obtaining more realistic dual radio models, but we also see the lack of inter-radio interference as an opportunity to evaluate the performance of low- and high-power radios separately in a controlled environment.

In our dual-radio system model when a local node discovers a remote node that has one or more content items of interest, the local node first turns on its high power radio interface. To ensure that the remote node also turns on its high power radio interface for uploading, we use a simple control protocol over the low power interface. The local node sends a *turn-on* control message to the remote node. When the remote node receives the message it turns on its high-power interface (if not already turned on). Afterwards it sends a *turned-on* message on the low-power interface to the local node. This message also contains the MAC address of the high power interface at the remote node and now the local node can request the discovered content items on the high power radio. When the local node is finished using the high-power radio it sends a *turn-off* message to the remote node to gracefully signal completion. Ungraceful session closure, such as when nodes move out of range or if the initiating node crashes, is handled by a soft-state timer.

## IV. RESULTS

### A. Measures and configurations

We investigate the performance both in terms of energy consumption and goodput (i.e. application throughput) from a system perspective. We study an open system and therefore it is important that the metrics are normalized with respect to the node sojourn time in the simulation. The system goodput is simply the sum of the number of bytes downloaded  $b_i$  by each node divided by the sum of the lifetimes of nodes in the simulation  $t_i$ , or  $G = \sum b_i / \sum t_i$ . We only count bytes of

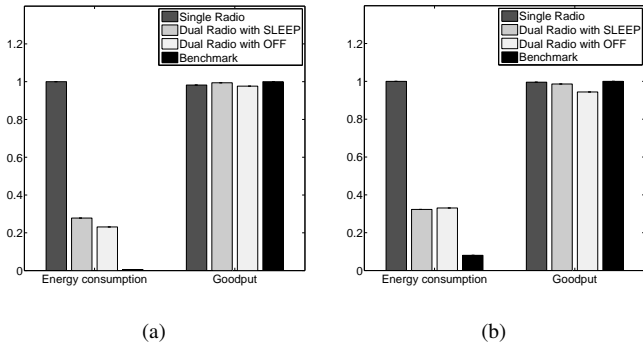


Fig. 3. Normalized energy consumption and normalized goodput for the Östermalm scenario (a) and the Subway scenario (b) with a 10m communication range.

fully downloaded content items, so the goodput is a measure of the system usefulness for the users, i.e. how much content it can provide. The energy consumption of the system is defined as the sum of the energy consumed by both the low- and the high-power radio normalized with respect to the sum of nodes' sojourn times. Thus, the energy consumption can be expressed as  $E = \sum (E_{lp_i} + E_{hp_i}) / \sum t_i$ .

In order to develop full understanding of the dual-radio system, in this work we examine the following configurations:

- 1) **Single radio** — an off-the-shelf solution as it is currently implemented in today's smartphones; a single radio interface is used for both service and content discovery as well as for actual content exchange.
- 2) **Dual radio with OFF** — a dual-radio system where the high-power radio is completely suspended when not in use; it is associated with a modest (2s) delay for turning on, but does not consume energy when suspended.
- 3) **Dual radio with SLEEP** — a dual-radio system where the high-power radio is put to sleep when not in use; the transition between sleeping and fully functional mode is instantaneous, but it comes at a price of a modest energy consumption (5% to 10% of the energy consumed when the high-power radio interface is turned on [8]).
- 4) **Benchmark** — an idealized system in which global knowledge is assumed for the location and the subscriptions of each node in the system; such a system concentrates on evaluation of the energy consumed for the actual data exchange and abstracts the discovery phase, thus it provides a lower bound for the energy consumption and an upper bound for the goodput.

In our model, if not stated otherwise, the primary data radio is based on the IEEE 802.11b MiXiM model and the secondary control radio is a simple low-power radio with IEEE 802.15.4-2006 beaconless MAC. The specifications of the low-power radio (such as range and energy profile) have been taken from the Texas Instruments ZigBee-Ready RF Transceiver (CC 2420). In the following evaluation, we have used the default parameters for both MAC layers, as stated in their corresponding specifications, although there have been propositions showing that tuning the IEEE 802.15.4 parameters might in certain cases lead to improved performance [18].

## B. 802.11b vs 802.15.4: A recap

Before delving into the evaluation results, we would like to make a small recap of the two MAC layer protocols we use in our dual radio system. Here we would like to stress the fact that although both protocols employ Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA), there are major design differences that affect the performance in an opportunistic context.

The IEEE 802.11b protocol is designed in a way that allows for faster adaptation to environmental changes. If the medium is sensed idle, the 802.11 protocol causes the node, wishing to transmit, to first defer for a fixed interval of  $50 \mu s$  (Distributed coordination function Inter-Frame Space, DIFS) and once this time elapses to check again the state of the medium. If it is still probed as free, a random back-off timer is started, during which the medium is constantly sensed. Such permanent observation of the environment allows for faster and more precise decisions to determine when a node can transmit. The back-off timer halts if another transmission is detected at any point in time. Moreover, once the channel is sensed free again for DIFS interval, the timer is resumed.

The beaconless 802.15.4-2006 protocol is more "careless" for environmental changes, because it was not initially designed for usage in dynamical environments. A transmission in 802.15.4 is possible only if the channel has been sensed free for the duration of the Clear Channel Assessment (CCA) after the expiration of the back-off period. The CCA is standardized to be 8 symbols ( $128 \mu s$ ) long. If the channel becomes busy during the CCA sensing, the back-off period is doubled, as opposed to the frozen back-off timer in 802.11. If a channel is not occupied by the end of the CCA, the transmitting station will enter transmitting mode (with a turnaround time of 12 symbols) and will then transmit the data in the medium. There is however no sensing during the turnaround interval, which is a potential period for collisions both with other data transmissions, as well as with acknowledgments.

## C. Initial content availability

In both scenarios we assume that all nodes carry devices and that there are 10 available feeds in the area, each feed containing 10 entries. In reality these nodes may as well represent only the proportion of nodes that carry devices in a denser environment. Every device is subscribed to one feed upon entering the observed area, and its content database is initially populated with 5 random entries (out of 10 entries available) on the feed. Thus, throughout its lifetime in the simulation, each node strives to obtain the rest of the entries that belong to its subscription feed. Entries have a mean size of 10 KB, and a standard deviation of 2 KB.

## D. Equal communication ranges

In our initial evaluation we have set the communication ranges of both radios to be equal. We are aware that 802.11 radios usually have a longer communication range than 802.15.4 radios. With a simple external power amplifier the range of an 802.15.4 radio can however easily be increased at the

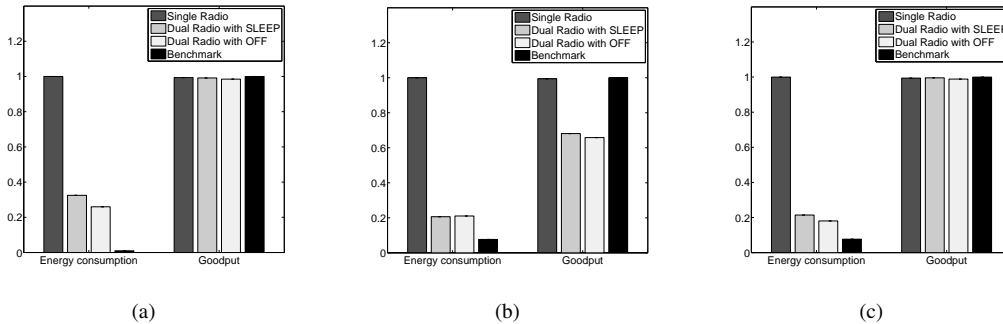


Fig. 4. Normalized energy consumption and normalized goodput for the Östermalm scenario (a) and the Subway scenario (b) with a 50 m communication range. (c) Normalized energy consumption and normalized goodput for the Subway scenario when the low-power radio adopts the 802.11b MAC.

expense of a slightly increased energy consumption and a lower bitrate [19]. Figures 3 and 4 show the performance of the Östermalm and the Subway scenario when the low-power and the high-power radio have equal communication ranges, respectively set to 10 m and 50 m. The energy consumption is normalized with respect to the single radio while the goodput is normalized with respect to the benchmark.

As we can see from Figure 3, both the Östermalm and the Subway scenario, although different in their nature, exhibit similar behavior when the high-power radio is suspended. For both cases the dual radio requires around 20% to 30% of the energy used with a single high-power radio. At the same time, we observe only a slight reduction in the amount of information obtained by the devices.

When nodes are in communication range, the delay associated with turning on the high-power interface may come at a cost of lost contacts. It is possible that by the time the high-power interface of a node is brought up, the peer carrying the content of interest has moved out of communication range and the opportunity for obtaining the data is lost. This is reflected in the slight performance drop in goodput when compared to the single radio and benchmark scenarios. We should however note that this drop may get larger when a mobility scenario of different character is considered.

To reduce the wake-up delay and minimize the probability of lost contact we also evaluate the effect of having a sleep state. The results show that the sleeping mode exhibits higher goodput values as compared to the complete high-power suspension while the difference in the amount of energy consumed in both modes is negligible.

Figure 4 presents the results of the same experiments conducted with a communication range of 50 m. We observe that while the performance metrics in the Östermalm scenario are just slightly decreased as compared to a system with a communication range of 10 m, the Subway scenario suffers serious losses in terms of goodput. This can be explained with the neighbor density in each of the scenarios and the performance of the 802.15.4 MAC under dense environmental conditions.

The increased range does not result in much higher neighbor densities for the Östermalm scenario, while it increases the neighbor densities dramatically for the more open spaces

of the Subway scenario. This affects the performance of the beaconless 802.15.4 MAC, because transmissions end up colliding, which exponentially increases the waiting times for next transmission attempts. Since we are examining a dynamic scenario where nodes do not experience long contact durations, the increased back-off durations result in partial node and content discovery as well as failures to wake-up the high-power radio for actual data transfers. As a verification of the poor performance of 802.15.4 MAC, we adopt the 802.11b MAC protocol for the low-power radio and perform the same experiment for the Subway scenario. In order to facilitate the comparison of the two scenarios, we assume that even though the low-power radio works with an 802.11b MAC, the energy consumption remains the same as for 802.15.4. The result is presented in Figure 4(c). A comparison between Figures 4(b) and 4(c) demonstrates that the 802.11b MAC obviates the 802.15.4 MAC performance issues. For an explanation of this phenomenon we can refer to the differences in the CSMA/CA mechanism of the two MAC protocols.

#### E. Different communication ranges

Power control of the low-power radio may offer a solution to the underperformance of the beaconless 802.15.4 MAC in dense scenarios. Hence, when in dense scenarios the range of the low-power radio can shrink in order to decrease the number of collisions and accommodate more possibilities for contacts, while in less populated areas it can expand and reach at longer distances. In order to check this idea, we run experiments both for the Östermalm and the Subway scenario where we set the range of the low-power radio at 10 m, while we allow the high-power radio to transmit at a 50 m range. The reason for not decreasing the high-power radio range is that we would like to allow the system to make use of contacts that occur when two nodes are on the periphery of the low-power radio range and walking away from each other. As long as they discover contents of interest on the low-power radio they can continue the download stage on the high-power radio, thus extending the range.

The results presented in Figure 5 illustrate that different ranges do contribute to the increase in goodput. These results should not however be considered in isolation, but should instead be compared to the corresponding cases for the Öster-




Fig. 5. Normalized energy consumption and normalized goodput for the Östermalm scenario (a) and the Subway scenario (b) with a 10 m communication range for the low-power radio and 50 m communication range for the high-power radio.

malm and the Subway scenario where both radios have a range of 50 m. In the Subway scenario (Figures 4(b) and 5(b)), one can see that decreasing the range of the low-power radio allows for increasing the goodput without significant changes in the energy consumption values. The energy spent on lost contacts in the case of equal ranges for both radios is now spent on actual data transmission and reception. The Östermalm scenario (Figures 4(a) and 5(a)) however presents us with a slightly deteriorated performance in terms of both goodput and energy consumption. These results suggest that an adaptive scheme can indeed lead to improving the performance of a dual-radio opportunistic system based on the environment where the system is operating.

## V. CONCLUSION

In this paper we have presented a simulation based evaluation of a dual-radio architecture in the context of opportunistic content distribution. The main benefit of a dual-radio system is that a simple low-power, low-bitrate radio can be used as a control channel to discover neighbors as well as contents, and only wake up the high-power, high-bitrate radio when necessary, thus potentially saving energy. To evaluate our system we use scenarios with different mobility characteristics.

Our results indicate that a dual radio system can significantly reduce the energy consumption of mobile devices operating in an opportunistic context, and these savings come at a price of a slight reduction in the goodput of the opportunistic network under low neighbor densities. Our initial measurements showed that smartphones exhaust their batteries four times faster when the 802.11 interface operates in ad-hoc mode. The proposed dual-radio architecture decreases the energy consumption to only 25% of that of a single radio system, thus recovering the operation time of the device. However, at high neighbor densities the results suggest that the current implementation of low-power chips (such as the ZigBee-like CC2420) may not be suitable for support of the low-power radio functionality in a dual-radio system, because of MAC layer design issues. We therefore recommend that dual-radio architectures should adopt an IEEE 802.11 MAC for the low-power radio if possible. Otherwise, we propose a

hybrid dual-radio solution, where the communication range of the low-power radio in a dual-radio system can adapt to the environment. The first results confirm that such adaptivity can contribute to improved performance of a dual-radio system in environments with high densities.

## REFERENCES

- [1] G. Karlsson, V. Lenders, and M. May, "Delay-Tolerant Broadcasting," *IEEE Transactions on Broadcasting*, vol. 51, no. 1, March 2007.
- [2] M. Papadopoulou and H. Schulzrinne, "Effects of power conservation, wireless coverage and cooperation on data dissemination among mobile devices," in *Proc. of ACM MobiHoc*, Long Beach, CA, USA, 2001.
- [3] V. Vukadinovic and G. Karlsson, "Spectral Efficiency of Mobility-Assisted Podcasting in Cellular Networks," in *Proc. ACM MobiOpp*, Feb. 2010.
- [4] M. Stemm and R. H. Katz, "Measuring and reducing energy consumption of network interfaces in hand-held devices," *IEICE Transactions on Communications, special Issue on Mobile Computing*, 1997.
- [5] O. Helgason, E. A. Yavuz, S. Kouyoumdjieva, L. Pajevic, and G. Karlsson, "A mobile peer-to-peer system for opportunistic content-centric networking," in *Proc. ACM SIGCOMM MobiHeld workshop*, 2010.
- [6] E. Shih, P. Bahl, and M. J. Sinclair, "Wake on wireless: an event driven energy saving strategy for battery operated devices," in *Proc ACM MobiCom*, 2002.
- [7] T. Pering, V. Raghunathan, and R. Want, "Exploiting radio hierarchies for power-efficient wireless device discovery and connection setup," in *Proc. IEEE VLSI*, 2005.
- [8] L. M. Feeney and M. Nilsson, "Investigating the energy consumption of a wireless network interface in an ad hoc networking environment," in *Proc. IEEE Infocom*, Apr. 2001.
- [9] L. M. Feeney and D. Willkomm, "Energy framework: an extensible framework for simulating battery consumption in wireless networks," in *Proc. SIMUTools*, 2010.
- [10] T. Pering, Y. Agarwal, R. Gupta, and R. Want, "Coolspots: reducing the power consumption of wireless mobile devices with multiple radio interfaces," in *Proc. MobiSys*, 2006.
- [11] H. Jun, M. H. Ammar, M. D. Corner, and E. W. Zegura, "Hierarchical power management in disruption tolerant networks using traffic-aware optimization," *Comput. Commun.*, vol. 32, no. 16, pp. 1710–1723, 2009.
- [12] "Legion Studio," <http://www.legion.com/>.
- [13] O. Helgason, S. T. Kouyoumdjieva, and G. Karlsson, "Does mobility matter?" in *Proc. IEEE/IFIP WONS*, 2010.
- [14] O. R. Helgason and K. V. Jónsson, "Opportunistic networking in OMNeT++," in *Proc. SIMUTools, OMNeT++ workshop*, 2008.
- [15] "MiXiM Framework," <http://mixim.sourceforge.net/>.
- [16] "OMNeT++," <http://www.omnetpp.org/>.
- [17] O. Helgason and S. T. Kouyoumdjieva, "Enabling multiple controllable radios in OMNeT++ nodes," in *Proc. ICST SIMUTools 2011, OMNeT++ workshop*, March, 2011.
- [18] D. Rohm, M. Goyal, H. Hosseini, A. Divjak, and Y. Bashir, "A simulation based analysis of the impact of IEEE 802.15.4 MAC parameters on the performance under different traffic loads," *Mob. Inf. Syst.*, vol. 5, pp. 81–99, January 2009.
- [19] D. Grini, "CC2420 with external PA," <http://focus.ti.com/lit/an/swra056b/swra056b.pdf>.