

Towards unified radio power management for wireless sensor networks

Guoliang Xing^{1*,†}, Mo Sha¹, Greg Hackmann², Kevin Klues^{2‡}, Octav Chipara² and Chenyang Lu²

¹*Department of Computer Science, City University of Hong Kong, Hong Kong*

²*Department of Computer Science and Engineering, Washington University in St. Louis, St. Louis, MO 63130, U.S.A.*

Summary

Many wireless sensor networks must sustain long lifetimes on limited energy resources. Two major approaches, transmission power control and sleep scheduling, have been proposed to reduce the radio power consumption in the transmission state and the idle state, respectively. In this paper, we first review existing transmission power control and sleep scheduling approaches and then describe a *Unified Radio Power Management* framework for the design and implementation of holistic radio power management solutions in wireless sensor networks. It has two key components: (1) a novel optimization approach called *Minimum Power Configuration* that minimizes the *aggregate* radio power consumption of all radio states and (2) a *Unified Power Management Architecture* (UPMA) that aims to support the flexible cross-layer integration of different power management strategies. A novel feature of UPMA is that it enables cross-layer coordination and joint optimization of different power management strategies that exist at multiple network layers. Copyright © 2008 John Wiley & Sons, Ltd.

KEY WORDS: Unified Power Management; transmission power control; sleep scheduling; network architecture

1. Introduction

Recent years have seen the deployments of wireless sensor networks (WSNs) in a variety of applications including habitat monitoring, border surveillance, and structural monitoring. WSNs in these applications must remain operational for long lifetimes on limited energy resources. For instance, due to the high cost for embedding sensors in bridges, a WSN deployed for structural health monitoring must continuously operate for years to be economically feasible. While energy harvesting techniques (e.g., solar panels [1]) can provide

additional energy at run time, the amount of energy available remains scarce. Therefore, power management is crucial for making WSNs viable in many real-world applications.

Radio is a major source of energy consumption in WSNs. Table I shows the power characteristics of two representative radio interfaces widely used in existing wireless sensor platforms. Two observations can be drawn from this table. First, the transmission power consumption has a wide tunable range, which offers opportunities for significant energy saving. Second, the power consumption in sleep state is several orders of

*Correspondence to: Guoliang Xing, Department of Computer Science, City University of Hong Kong, Hong Kong.

†E-mail: glxing@cityu.edu.hk

‡Now with Stanford University.

Table I. The power consumption of two radio platforms in wireless sensor networks [2].

Platforms	Transmission (mW)	Reception/Idle (mW)	Sleep (μ W)
Chipcon CC1000	19.8–80.1	22.2	0.6
Chipcon CC2420	25.5–52.2	59.1	60

magnitude lower than in other states. In accordance with these observations, two major approaches have been proposed to achieve power-efficient communication in WSNs. *Transmission power control* aims to reduce the overall transmission power of a network by adjusting the transmission power at each node. *Sleep scheduling* reduces the radio energy wasted in idle state by turning off radios when not in use.

Although a multitude of power management strategies have been developed for WSNs, WSNs still face two important challenges in radio power management. First, existing power management approaches aim at reducing the power consumed in only *one* radio state. As a result, they become ineffective under different network conditions. Transmission power control only reduces the transmission power, and hence is not effective for applications where nodes do not transmit frequently. Sleep scheduling, on the other hand, only reduces the power consumption in the idle state, and thus is not effective for heavy-load applications where power consumption is dominated by packet transmissions. Therefore, in order to improve the lifetime of a WSN in different applications, both transmission power control and sleep scheduling must be employed in a coordinated fashion. However, to achieve this, two important challenges must be overcome. First, power management strategies are commonly developed in isolation. Consequently, when multiple power management strategies are used simultaneously, their decisions may be contradictory. Mechanisms for resolving such conflicting decisions must be employed. Second, sleep scheduling policies are commonly implemented as part of the medium access control (MAC) layer, while transmission power control policies are implemented at the network layer. The effective coordination between the two strategies is hindered by the strict layering of the network stack and the monolithic implementations of each layer. Therefore, a new architecture for the network stack that allows for coordinating power management strategies employed at different layers is necessary.

In this paper, we first provide a brief survey of existing transmission power control and sleep scheduling

schemes. We then present a *unified* approach to radio power management, which aims to integrate the existing power management strategies. To this end, we first describe a novel optimization approach called *Minimum Power Configuration* (MPC). In contrast to the existing power management algorithms that treat different radio states in isolation, MPC provides a unified optimization approach that minimizes the total energy consumption in all radio states based on network workloads. We then describe a *Unified Power Management Architecture* (UPMA) for flexible integration of different power management strategies. UPMA facilitates cross-layer coordination of different power management strategies and defines standardized interfaces between power management strategies and the rest of the network protocol stack.

The rest of article is structured as follows. Sections 2 and 3 review existing transmission power control and sleep scheduling approaches, respectively. We then present the Minimum Power Configuration (MPC) approach (Section 4) and the Unified Power Management Architecture (Section 5). We conclude the paper in Section 6.

2. Transmission Power Control

Transmission power control adjusts the radio transmission power in order to achieve lower power consumption, while preserving certain network properties. Generally, these schemes reduce the power to the minimum level required to achieve a “good” topology. The characteristics that define a topology’s quality vary from scheme to scheme; commonly targeted properties include bounds on the network connectivity, average node degree, and packet reception ratio (PRR). Such schemes primarily achieve energy savings by directly reducing the cost of transmitting a packet over the radio. Moreover, reducing the transmission power lowers the number of packet retransmissions by reducing contention among nearby nodes.

Traditionally, transmission power control schemes have assumed that each node has a circular radio range. A comprehensive survey of such schemes can be found in [3]. However, empirical studies have shown that the circular radio model used by many of these schemes is unrealistic [4]. First, lossy and asymmetric links are prevalent in WSNs. Zhao *et al.* [4] reported that a third of the links in a test-bed composed of 60 Mica motes experienced more than 30% packet loss even under light workloads. Consequently, up to 80% of the total energy consumption of the radio was attributed to packet loss

[4]. Moreover, the wireless communication between nodes in WSNs suffers from significant variations with time and environments, and hence the packet reception performance of links is stochastic in nature [5]. Henceforth, we focus on several representative transmission power control schemes designed based on the realistic properties of WSNs.

LMA and LMN [6] periodically adjust a node's transmission power to ensure the number of one-hop and two-hop neighbors (respectively) fall within some bounds. These algorithms have a very low memory footprint and can react to dynamic changes in link conditions. However, LMA and LMN may partition the network into disconnected clusters, since they do not consider global network connectivity when selecting neighbors.

XTC [7] removes "bad" links from the network graph G through a series of local decisions. A node u establishes a local order \prec_u over its one-hop neighbors using some link quality metric. Nodes exchange these orderings \prec with their one-hop neighbors and process these exchanged orderings in order to filter out inefficient links. Specifically, a node u excludes a neighbor v from its neighbor table if these two nodes share a common neighbor w , and if the link from v to w is of higher quality than the direct link from v to u . The authors prove that XTC always preserves network connectivity, and demonstrate that XTC significantly reduces the average-case node degree in practice. They suggest that nodes order their neighbors based on the minimum power level required to establish a connection, so that XTC will generate a topology consisting mainly of low-power links. However, the authors' discussion focuses on the connectivity of the produced topologies, and does not include experimental data to show the overall effect of XTC on power consumption.

ATPC [8] generates a runtime model that correlates transmission power and link quality for each pair of neighbors. Using this model, each node adjusts its transmission power level on a per-packet basis in order to achieve some bound on link quality. ATPC also incorporates a closed feedback loop that dynamically adjusts the link quality model as network conditions change. Testbed experiments demonstrate that ATPC consistently achieves high PRR and reduced transmission power, even when subjected to environmental changes. However, ATPC's model is based in part on extensive experimental data which must be collected *a priori* for each location.

PCBL [9] collects link statistics for each one-hop neighbor at each power level supported by the radio.

Transmission power is selected on a per-link basis, using the lowest power level that achieves some upper bound on PRR whenever possible. Links which fall below some lower bound on PRR at all power levels are never used for transmissions. PCBL has been shown to significantly reduce power consumption in a real WSN testbed while still achieving good PRR. Because PCBL creates a static neighbor table at each power level at boot time, it has a high bootstrapping cost and limited robustness in dynamic settings.

While ATPC and PCBL are designed to maintain desired qualities for individual links, multi-hop WSNs may require certain end-to-end quality of routes comprising multiple hops. CTC [10] aims to replace high-power paths in multi-hop wireless sensor networks with low-power alternatives. Nodes collect the minimum transmission count for their neighbors at all power levels and propagate this information among their two-hop neighborhood. Each node independently replaces max-power links with lower-power paths, subject to a user-provided dilation of transmission count (DTC) bound. For example, a max-power link (u, v) is replaced with a lower-power path (u, w, v) only if $\frac{\Gamma(u,w)+\Gamma(w,v)}{\Gamma(u,v)} \leq t$ for some DTC bound t , where $\Gamma(x, y)$ is the expected number of transmissions of the link from x to y . Each node additionally simulates the execution of CTC on its one-hop neighbors to ensure that path-replacement decisions are consistent across the entire network. This approach guarantees a lower bound on global link quality, even though CTC is entirely distributed. A disadvantage of CTC is that it incurs a high bootstrapping cost as the PRR of links at different transmission power levels need to be periodically collected.

3. Sleep Scheduling

As shown in Table I, the idle listening state of radio consumes a non-negligible amount of power compared to the transmission state. Sleep scheduling aims to reduce radio energy consumption by turning radios off when not in use. There are two basic approaches, namely *duty-cycle* and *backbone* based sleep scheduling. A backbone-based scheme aims to reduce the spatial density of active nodes while maintaining a desirable network topology. A comprehensive survey on existing backbone-based sleep scheduling schemes can be found in [11]. In this section, we focus on the duty-cycle-based sleep scheduling schemes.

Under duty-cycle sleep scheduling, nodes switch their radios between an active state and sleep state to reduce the idle listening time. Nodes only communicate

when in the active state. The duty cycles of different nodes can be synchronous [12–14] or asynchronous [15–18]. In synchronous duty cycling [12,19], nodes agree on a schedule that specifies when radios are active and asleep in a period. Such an approach is adopted by several MAC protocols. A well-known example is S-MAC [19]. To reduce the synchronization overhead, S-MAC introduces a loose synchronization mechanism called virtual clustering. Nodes with the same schedule form a virtual cluster and the nodes that lie on the border of two virtual clusters follow the schedules of both clusters, which maintains connectivity across the network. To reduce the energy consumed in overhearing unnecessary traffic, S-MAC uses an in-channel signaling mechanism that allows a node to turn off its radio when it is not the receiver of the on-going traffic. A drawback of the synchronous duty cycling is the increased communication delay because packets can only be transmitted when all nodes on the communication path are active. To achieve a better trade-off between power consumption and communication delay, several adaptive listening mechanisms have been proposed [14,19,20]. With this mechanism, a node remains active even after its active interval in the sleep schedule when it may be the current or next-hop receiver of on-going packet transmissions.

Different from synchronous duty cycling in which nodes maintain the same sleep schedule, asynchronous duty cycling allows each node to have a different sleep schedule. In B-MAC [16], nodes in the networks schedule their duty cycles independently and employ a preamble-based mechanism to synchronize the communication between two nodes. If a node has data to send, it first transmits a preamble composed of a regular bit pattern “010101 . . .”. To guarantee that the receiver will hear the preamble and turn on its radio for incoming packets, the preamble must last longer than the sleep period. Due to its on-demand wakeup mechanism, B-MAC incurs low overhead as nodes do not need to periodically synchronize their sleep schedules. However, the nodes within one hop from the sender will be woken up by the preamble, which causes unnecessary overhead. X-MAC [21] addresses this issue by embedding the receiver’s address in the preamble. Moreover, X-MAC reduces overhead of preamble by transmitting a series of short preamble packets. The receiver then responds immediately after hearing one preamble packet, which shortens the length of preamble as well as reduces the communication delay due to sleep scheduling.

Time division multiple access (TDMA) is another asynchronous duty cycling technique in which nodes

wake up at different scheduled time slots. In TDMA-based protocols, nodes transmit only during their own time slots, and hence can turn off their radios in other time slots. Several different TDMA-based protocols have been proposed for use in WSNs, such as TRAMA [17], DRAND [22], and GTS portion of 802.15.4 [23]. One limitation of TDMA-based sleep scheduling protocols is that nodes’ schedules can be very sensitive to changes in network traffic or network topology, and all nodes sharing a schedule must remain synchronized with one another.

Recently, several hybrid TDMA protocols, such as SCP [24], Z-MAC [18], and Funneling MAC [25] have been proposed. SCP combines scheduled contention and channel polling by synchronizing the time that nodes wake up to sample the radio channel. This synchronization allows sender nodes to send data with very short preambles. Z-MAC employs a TDMA-style slot allocation for all nodes, but allows nodes to contend for access to other nodes’ slots using channel polling. This approach combines TDMA’s low channel contention with channel polling’s high throughput. Finally, nodes equipped with Funneling MAC contend for channel access in the majority of the network via CSMA/CA, while using TDMA in regions close to sink nodes where nodes experience high contention. Funneling MAC alleviates contention in the most active areas of the network without requiring other nodes to create and maintain TDMA schedules.

Although duty cycling reduces the idle listening power of radios, it introduces communication delay because nodes can only transmit/receive packets when they are active. A technique known as Wake-On-Radio [26–28] has been proposed to address this issue. Specifically, a second low-power radio is employed to listen on the channel and wake up the primary radio to receive incoming packets.

4. Minimum Power Configuration

In order to maximize the system life of a WSN, the energy consumption in all radio states must be minimized. However, the existing approaches only optimize the energy consumption of a particular radio state resulting in unnecessary power waste. Transmission power control reduces the transmission energy of wireless nodes but does not consider the idle energy. Sleep scheduling can reduce the idle energy by scheduling idle nodes to sleep but does not optimize the transmission energy. In this section, we discuss a new approach called the *Minimum*




Fig. 1. (a) Two communication paths from a to c: $a \rightarrow c$ or $a \rightarrow b \rightarrow c$. (b) The power consumption of two different paths versus the data rate.

Power Management [29,30] that minimizes the total energy consumption of all radio states by integrating transmission power control and sleep scheduling.

4.1. An Illustrative Example

We now illustrate the basic idea of our approach with a simple example. Figure 1 shows three nodes a , b , and c with the same bandwidth of B bps. a needs to send data to c at the rate of R bps. There are two network configurations to accomplish the communication between a and c : (1) a communicates with c directly using transmission range $|ac|$ while b switches to sleep or (2) a communicates with b using transmission range $|ab|$ and b relays the data from a to c using transmission range $|bc|$. Minimizing the total energy of all nodes is equivalent to minimizing the average power consumption of the three nodes in all radio states. The average power consumption of each node can be computed by the following equation:

$$\bar{P} = \sum_{i=\{\text{Idle}, \text{Tx}, \text{Rx}, \text{Sleep}\}} \text{fraction of time in state } i \times \text{power in state } i$$

We denote the average power consumption under the two configurations as P_1 and P_2 , respectively. P_1 and P_2 can be computed as follows:

$$P_1 = \frac{R}{B} \cdot P_{tx}(|ac|) + \frac{R}{B} \cdot P_{rx} + 2 \left(1 - \frac{R}{B}\right) \cdot P_{id} + P_s \quad (1)$$

$$P_2 = \frac{R}{B} \cdot (P_{tx}(|ab|) + P_{tx}(|bc|)) + \frac{2R}{B} \cdot P_{rx} + \left(3 - \frac{4R}{B}\right) \cdot P_{id} \quad (2)$$

where $P_{tx}(d)$, P_{rx} , P_{id} , and P_s represent the radio power consumption in transmission (with a range of distance d), reception, idle, and sleeping, respectively. Each term in P_1 or P_2 is the product of power consumption in a radio state and the fraction of time the radio operates in that state. For example, in the first term of P_2 , $P_{tx}(|ab|) + P_{tx}(|bc|)$ is the transmission power of nodes a and b , and $\frac{R}{B}$ is the fraction of time nodes a and b operate in the transmission state. Similarly, the second term of P_2 represents the contribution of the reception power of nodes b and c . In the third term of P_2 , P_{id} is the idle power, and $3 - \frac{4R}{B}$ is the sum of the fractions of time when nodes a , b , and c stay in the idle state. Specifically, node a is idle $1 - \frac{R}{B}$ of the time because it becomes idle when not transmitting to b , node b is idle $1 - \frac{2R}{B}$ of the time because it becomes idle only when neither transmitting to c nor receiving from a , and node c is idle $1 - \frac{R}{B}$ of the time because it becomes idle when not receiving from b .

For the given radio parameters and node locations, all symbols except R are constant in the expressions of P_1 and P_2 . We plot P_1 and P_2 in Figure 1 under a possible setting of radio parameters and node locations. We can see that $P_1 > P_2$ when the data rate exceeds a threshold R_0 . To get a concrete estimation on R_0 , we now apply the parameters of the CC1000 radio on Mica2 motes to the above example. For a 433 MHz CC1000 radio, the bandwidth is 38.4 kbps. There are a total of 31 transmission power levels, each of which leads to a different transmission range. Suppose $P_{tx}(|ac|)$ is equal to the maximum transmission power 80.1 mW. $P_{tx}(|ab|)$ and $P_{tx}(|bc|)$ are equal to the medium transmission power 24.6 mW. P_{id} , P_{rx} , and P_s are 24 mW, 24 mW, and 6 μ W, respectively. Using this information, it can be calculated that relaying through node b is more power efficient when the data rate is above 16.8 kbps.

This example illustrates the basic idea of MPC. (1) When network workload is low, energy consumption of a network is dominated by the idle state of the radio. In such a case, scheduling nodes to sleep saves the most energy. MPC therefore uses high power communication links between nodes which allows any nodes that would otherwise be used as relays to sleep. (2) When network workload is high, the transmission energy dominates the total energy consumption of a network. Since transmission power increases quickly with distance, MPC uses shorter communication ranges and relays data through multiple nodes to save energy. Figure 2 shows two different network topologies configured by MPC for the high and low workload scenarios, respectively.

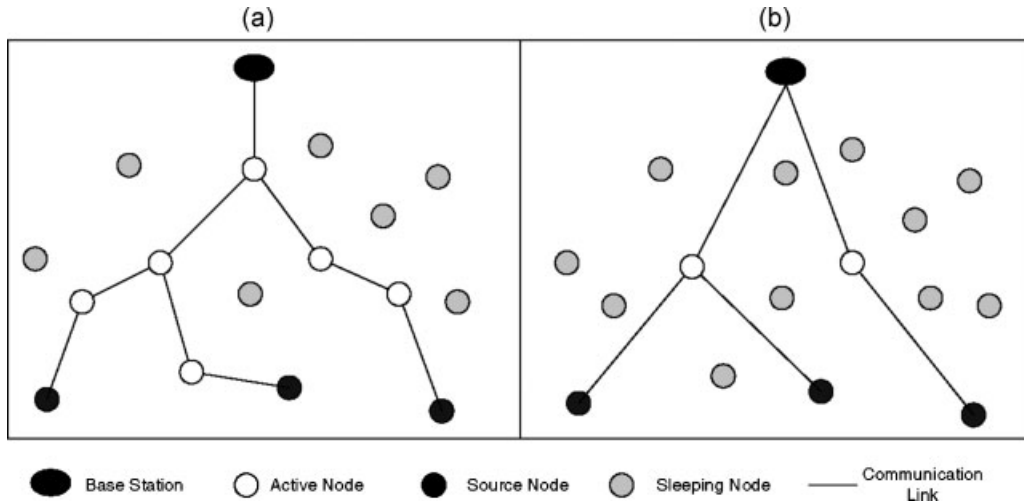


Fig. 2. Network topologies configured by MPC for high and low workload scenarios. (a) Data rates of sources are high. MPC chooses six active nodes and uses low transmission range; (b) Data rates of sources are low. MPC chooses two active nodes and uses high transmission range.

4.2. Minimum Power Configuration in Data Collection

In this section, we discuss how to apply the idea of MPC to reducing the energy consumption in data collection, which is a fundamental communication paradigm in WSNs. In data collection, a set of source nodes periodically sample the environments and send the data to the base station. The data rate of a source is often determined by the type of the data it samples. For example, sensor nodes in a typical habitat monitoring application report temperature measurements of petrel nests every 5–10 minutes in order to provide enough data for studying petrels' nesting behavior. Moreover, different nodes may serve as source nodes when triggered by the occurrence of interesting events.

We now describe an online algorithm called *Incremental Shortest-path Tree Heuristic (ISTH)* [29] for the minimum power configuration of data collection. ISTH executes in an online fashion as a node may start sending data to the base station at any time when triggered by an event of interest. ISTH finds a route from each source node to the base station and activates all the nodes on the route. Other nodes are scheduled to operate with low duty cycles. The choices of the active nodes on the route and their transmission power are jointly determined based on the data rate of the source. As the low duty nodes sleep in most of time, we focus on reducing the total energy consumption of active nodes.

As shown by the example in Section 4 (see (2)), the average power consumption of a node u , $P(u)$, in a

multi-hop path with data rate R_i can be computed as follows:

$$\begin{aligned} P(u) &= \frac{R_i}{B} \cdot (P_{tx}(u, v) + P_{rx}) + \left(1 - \frac{2R_i}{B}\right) \cdot P_{id} \\ &= \frac{R_i}{B} \cdot (P_{tx}(u, v) + P_{rx} - 2P_{id}) + P_{id} \\ &= R_i \cdot C_{u,v} + P_{id} \end{aligned} \quad (3)$$

where node v is the next-hop node of u on the path and $C_{u,v} = (P_{tx}(u, v) + P_{rx} - 2P_{id})/B$. From (3), we can see that the average power of each active node is equal to the sum of P_{id} and $C_{u,v}$ for each unit of data that flows through it. We now explain the basic idea of ISTH; a detailed description can be found in [29]. For each source with data rate r_i , ISTH finds the shortest route to the base station according to the following cost metric:

$$c(u, v) = \begin{cases} \frac{R_i}{B} \cdot C_{u,v} + P_{id} & \text{u is sleeping} \\ \frac{R_i}{B} \cdot C_{u,v} & \text{u is active} \end{cases} \quad (4)$$

This cost metric is designed based on two principles: (1) If a node is not on the route from any source to the base station, the cost of including it on a new route is $\frac{R_i}{B} \cdot C_{u,v} + P_{id}$. This cost accounts for the total power consumption in all radio states as shown in (3). Moreover, it allows for ISTH to choose power-efficient routes adaptively according to the data rate. When r_i is high, transmission power dominates the total network power consumption, and hence the routes that have a

small number of high-power links are more expensive under function (4). Consequently, ISTH chooses long routes with more low-transmission power links. On the other hand, when r_i is low, idle listening power dominates the total network power consumption. ISTH then chooses long routes with more low-transmission links to reduce the idle listening power. (2) If a node is already on a route from a source to the base station, the cost of including it on a new route is $\frac{R_i}{B} \cdot C_{u,v}$, which does not include the idle power P_{id} . This is because P_{id} has been counted by the existing routes, as each node only incurs one P_{id} (according to (3)) independent of the amount of data that are transmitted through it.

To account for the additional energy cost of re-transmissions caused by packet loss, the cost function defined in (3) can be revised as follows. Let $PRR(u, v, P_{tx})$ represent the PRR when u communicates with v using transmission power P_{tx} . Note that $PRR(u, v, P_{tx})$ depends on the quality of both forward and reverse links between u and v when an automatic repeat request (ARQ) scheme is used. The expected transmission power cost when u communicates with v with P_{tx} on the lossy links can be estimated as $P_{tx}/PRR(u, v, P_{tx})$.

Simulation results [29] show that MPC significantly outperforms two other data collection approaches named Minimum Transmission (MT) and Minimum Transmission Power (MTP). MT is shown to be more reliable than the hop-count based routing scheme when given lossy networks [31]. A node in MT chooses the next hop node with the minimum expected number of transmissions to the sink. A node in MTP chooses the next hop node with minimum total expected transmission power to the sink. Except for the consideration for unreliable links, MTP is similar to the minimum power routing schemes studied in [32,33]. It is shown in [29] that MPC reduces the network energy consumption of 30% over MTP and 26% over MT.

5. Unified Power Management Architecture

The MPC approach discussed in Section 4 demonstrates the significant advantage of integrating power management strategies at different layers. However, implementing integrated cross-layer power management solutions is challenging on existing network platforms, as the existing power management strategies are usually tightly coupled with network protocols and other system functionality. This monolithic approach

has led to standalone solutions that cannot be easily reused or extended to other applications or platforms. Moreover, different power management strategies make different and sometimes even conflicting assumptions about the rest of the system with which they need to interact. The lack of architectural support has made it difficult to integrate existing power management protocols within a diverse set of applications and network platforms. We now present the Unified Power Management Architecture (UPMA) [34], which supports flexible integration of power management strategies in multiple layers.

5.1. Motivation for UPMA

Current power management strategies often adopt monolithic implementations in which power management is tightly coupled with a particular network protocol stack. As a result, a system is often limited to specific power management strategies that cannot be easily extended or replaced. For example, sleep scheduling is often implemented as part of MAC protocols, while power control is often integrated with routing or topology maintenance protocols. Furthermore, many routing protocols use specific power control schemes to compute a set of routing metrics based on transmission power. It would be more flexible to separate the power control functionality from any specific routing protocols that implement them, and simply provide an interface to fetch the values of any cost metrics. Higher level services like TinyDB [35] also employ multiple power management strategies that stretch across multiple layers and are specifically designed to work together. Although each of these monolithic approaches are often slightly more computationally efficient, they have largely impeded the interoperability of different power management protocols, and the overall synergy between different research efforts.

While current research efforts mainly focus on the use of a single power management protocol, a unified architecture is needed to effectively compose different protocols together to form a single coherent power management solution. Each individual protocol may be sub-optimal since it only reduces the energy consumption in some subset of its radio states. Power control only reduces the transmission power of nodes, while sleep scheduling reduces the idle power. Furthermore, some sleep scheduling protocols [36,37] have been specifically designed for data collection applications that impose periodic low network traffic, while backbone maintenance protocols are designed for applications (e.g., real-time detection and tracking) in which

message delivery latency is extremely important. In order to minimize the total energy consumption of a network, application developers must effectively integrate the use of different power management protocols across different layers. The MPC approach shows that the optimal integration of different power management protocols requires careful cross-layer consideration of the radio characteristics, routing choices, and network workload imposed by the application. The existence of a unified architecture within which these tasks could be performed would be very beneficial.

5.2. The Design of UPMA

The UPMA is based on a set of architectural abstractions that are designed with the following principles in mind. (1) They should support the development of a multitude of different power management protocols, each having their own independent implementations. (2) They should contain a set of standardized interfaces between all power management protocols and any other components in the system. (3) They should allow components to be integrated into the architecture that are capable of performing cross-layer coordination between power management protocols existing at different layers.

As shown in Figure 3, the design of UPMA aims to meet each of these requirements. (1) Power management protocols exist as independent entities at both the network layer as well as the data link layer. (2) Communication takes place between these protocols and other components in the system through a standard

set of interfaces. (3) Cross-layer coordination can be achieved through the proper implementation of different *Power Coordination Tables* and their corresponding *Power Coordinator* component.

The standard set of interfaces encapsulate the requirements of the representative power management protocols. Sleep scheduling protocols need to be able to (1) turn the radio on and off (*PowerControl*), (2) perform clear channel assessment on the radio channel (*ChannelMonitor*), and (3) set the preamble length of an outgoing packet (*PreambleLength*). Power control protocols need to (1) set the transmission power level that a packet should be transmitted at (*TxPower*), and (2) specify the routing cost for use by network protocols existing in the system (*Cost*).

The Power Coordinator and its corresponding Power Coordination Tables are configured differently based on which power management protocols are being used in the system. These components can be instantiated as necessary to meet the power constraints of any applications running on top of them. As a simple example, consider two applications specifying two different duty cycles for a single underlying sleep scheduling protocol. The Power Coordination Tables store the on and off times required for each duty cycle, and the Power Coordinator combines these values to produce a sleep schedule satisfying the on time requirements of both. A comprehensive evaluation of such a configuration has been performed in [34], with results indicating that flexibility is indeed increased without incurring a significant performance penalty.

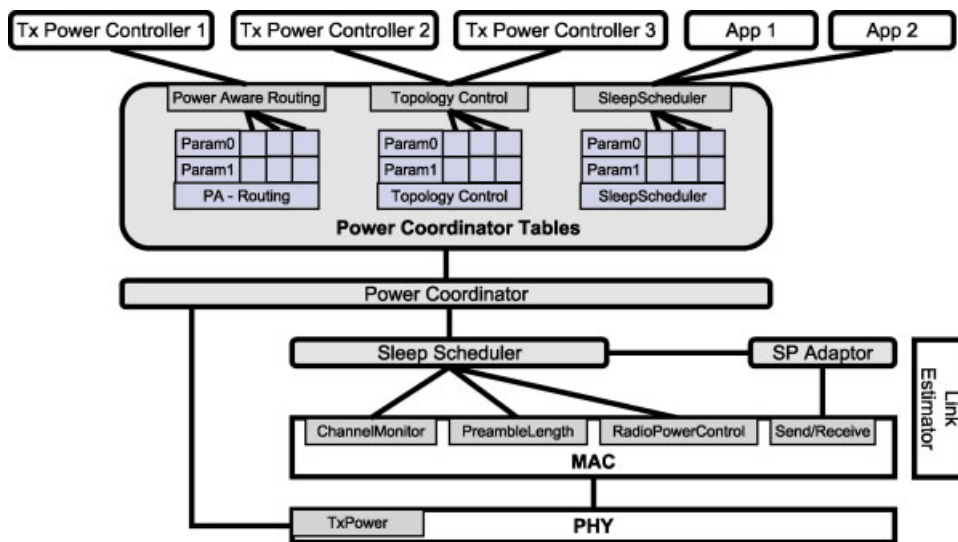


Fig. 3. The proposed Unified Radio Power Management Architecture.

5.3. Proposed Implementation of MPC in UPMA

We now briefly discuss how the MPC approach can be implemented in UPMA. The implementation of MPC can be broken down into two interleaving components: a sleep scheduling component and a power-aware routing component. When a node starts routing a data flow, the sleep scheduler stops duty-cycling the node. This state transition triggers the power-aware routing component to optimize the transmission power of the node and to assign the node a lower routing cost. As a result, data tend to be routed through only those nodes that are currently active, resulting in less energy wasted by idle listening. Within UPMA, MPC can be realized by creating appropriate power control and sleep scheduling components, and implementing a cross-layer optimization protocol within the Power Coordinator component. Once the node starts routing a data flow, the Power Coordinator can modify the values in the Power Coordination Tables to indicate that a node should always be powered on. At the same time, it can compute any new routing costs based on this change and update this value in the appropriate network protocol through a standard interface.

5.4. Current Implementation of UPMA

We have implemented the following components of UPMA in TinyOS 2.x, the second generation of the widely adopted standard operating system for WSNs: (1) the interfaces between the MAC layer and sleep scheduling algorithms [34]; (2) the framework for integrating different sleep scheduling strategies [34]; (3) a component-based MAC Layer Architecture (MLA) architecture for implementing different power-efficient MAC protocols [38]; (4) five different MAC protocols within MLA [38]. Microbenchmark results [34] demonstrate that UPMA incurs a negligible decrease in performance when compared to existing monolithic implementations of MAC layer power management strategies. Several case studies show that the power management requirements of multiple applications can be easily coordinated within UPMA, sometimes even resulting in better power savings than any one of them can achieve individually. On-going work includes supporting upper layer protocols such as topology control and cross-layer coordination. More information about the UPMA project can be found at <http://www.cse.wustl.edu/~lu/upma.html>.

6. Conclusion

Power management is a fundamental challenge in WSNs. Previous work has focused on different strategies for minimizing the power consumption in individual ratio states. In this paper, we first review two major power management approaches, transmission power control and sleep scheduling. We then discuss an emerging approach called *unified radio power management* composed of two key components. Minimum Power Configuration (MPC) is a novel optimization approach that minimizes the total power consumption of all ratio states by integrating transmission power control and sleep scheduling into a unified framework. The transmission power choices and sleep scheduling decisions of nodes are coordinated according to the current network workload. The Unified Power Management Architecture (UPMA) supports the flexible integration of different power management strategies. A key feature of UPMA is that it enables cross-layer coordination and joint optimization of different power management strategies that exist at multiple network layers while allowing them to have independent implementations.

Acknowledgements


The work described in this paper was partially supported by the Research Grants Council of Hong Kong under grants RGC 9041266 and the NSF under a NeTS-NOSS grant no CNS-0627126.

References

1. Kansal A, Hsu J, Zahedi S, Srivastava MB. Power management in energy harvesting sensor networks. *Transaction on Embedded Computing Systems* 2007; **6**(4): 120–158.
2. Crossbow. Mica and mica2 wireless measurement system datasheets. 2003.
3. Santi P. Topology control in wireless ad hoc and sensor networks. *Technical Report IIT-TR-04*, Istituto di Informatica e Telematica, Pisa, Italy, 2003.
4. Zhao J, Govindan R. Understanding packet delivery performance in dense wireless sensor networks. In *Sensys*, Los Angeles, CA, November 2003.
5. Cerpa A, Wong JL, Kuang L, Potkonjak M, Estrin D. Statistical model of lossy links in wireless sensor networks. In *IPSN*, 2005.
6. Kubisch M, Karl H, Wolisz A, Zhong LC, Rabaey J. Distributed algorithms for transmission power control in wireless sensor networks. *IEEE Wireless Communications and Networking* 2003; **1**: 558–563.
7. Wattenhofer R, Zollinger A. Xtc: A practical topology control algorithm for ad-hoc networks. In *WMAN*, 2004.
8. Lin S, Zhang J, Zhou G, Gu L, Stankovic JA, He T. Atpc: adaptive transmission power control for wireless sensor networks. In *SenSys*, 2006.

9. Son D, Krishnamachari B, Heidemann J. Experimental study of transmission power control and blacklisting based link quality control in wireless sensor networks. *Technical Report ISI-TR-629*, Information Sciences Institute, 2007.
10. Xing G, Lu C, Pless R. Localized and configurable topology control in lossy wireless sensor networks. In *16th International Conference on Computer Communications and Networks (ICCCN)*, 2007.
11. Wang L, Xiao Y. A survey of energy-efficient scheduling mechanisms in sensor networks. *Mobile Network and Applications (MONET)* 2006; **11**(5): 723–740.
12. IEEE. Wireless lan medium access control (mac) and physical layer (phy) specifications. In *IEEE Standard 802.11*, 1999.
13. Ye W, Heidemann J, Estrin D. An energy-efficient mac protocol for wireless sensor networks. In *INFOCOM*, 2002.
14. van Dam T, Langendoen K. An adaptive energy-efficient mac protocol for wireless sensor networks. In *Proceedings of the First ACM Conference on Embedded Networked Sensor Systems (SenSys)*, Los Angeles, CA, USA, 2003.
15. Zheng R, Hou JC, Sha L. Asynchronous wakeup for ad hoc networks. In *Proceedings of the 4th ACM International Symposium on Mobile Ad Hoc networking and Computing (MobiHoc)*, Annapolis, MD, USA, 2003.
16. Polastre J, Hill J, Culler D. Versatile low power media access for wireless sensor networks. In *SenSys*, 2004.
17. Rajendran V, Obraczka K, Garcia-Luna-Aceves J. Energy-efficient, collision-free medium access control for wireless sensor networks. In *Proceedings of the First ACM Conference on Embedded Networked Sensor Systems (SenSys)*, Los Angeles, CA, USA, 2003.
18. Rhee I, Warrier A, Aia M, Min J. Z-mac: a hybrid mac for wireless sensor networks. In *Proceedings of the Third ACM Conference on Embedded Networked Sensor Systems (SenSys)*, San Diego, CA, USA, 2005.
19. Ye W, Heidemann J, Estrin D. Medium access control with coordinated, adaptive sleeping for wireless sensor networks. *IEEE/ACM Transactions on Networking*, June 2004; **12**(3): 493–506.
20. Zheng R, Kravets R. On-demand power management for ad hoc networks. In *INFOCOM*, 2003.
21. Buettner M, Yee GV, Anderson E, Han R. X-mac: a short preamble mac protocol for duty-cycled wireless sensor networks. In *SenSys*, 2006.
22. Rhee I, Warrier A, Min J, Xu L. Drand: distributed randomized tdma scheduling for wireless ad-hoc networks. In *MobiHoc*, 2006.
23. IEEE. Wireless medium access control (mac) and physical layer (phy) specifications for low-rate wireless personal area networks (lr-wpans). *IEEE Standard 802.15.4*, 2003.
24. Ye W, Silva F, Heidemann J. Ultra-low duty cycle mac with scheduled channel polling. In *SenSys*, 2006.
25. Ahn G-S, Hong SG, Miluzzo E, Campbell AT, Cuomo F. Funneling-mac: a localized, sink-oriented mac for boosting fidelity in sensor networks. In *SenSys*, 2006.
26. Shih E, Bahl P, Sinclair MJ. Wake on wireless: an event driven energy saving strategy for battery operated devices. In *Proceedings of the Eighth Annual International Conference on Mobile Computing and Networking*, Atlanta, GA, USA, 2002.
27. Miller MJ, Vaidya NH. A MAC protocol to reduce sensor network energy consumption using a wakeup radio. *IEEE Transactions on Mobile Computing*, 2005; **4**(3): 228–242.
28. Gu L, Stankovic JA. Radio-triggered wake-up capability for sensor networks. In *Proceedings of the 10th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, Toronto, Canada, May 2004.
29. Xing G, Lu C, Zhang Y, Huang Q, Pless R. Minimum power configuration in wireless sensor networks. In *Proceedings of the 6th ACM international symposium on Mobile ad hoc networking and computing (MobiHoc)*, 2005.
30. Xing G, Lu C, Zhang Y, Huang Q, Pless R. Minimum power configuration for wireless communication in sensor networks. *ACM Transactions Sensor Networks* 2007; **3**(2): 200–233.
31. Woo A, Tong T, Culler D. Taming the underlying challenges of reliable multihop routing in sensor networks. In *SenSys*, 2003.
32. Doshi S, Bhandare S, Brown TX. An on-demand minimum energy routing protocol for a wireless ad hoc network. *SIGMOBILE Mobile Computing Communications Review* 2002; **6**(3): 50–66.
33. Doshi S, Brown TX. Minimum energy routing schemes for a wireless ad hoc network. In *INFOCOM*, 2002.
34. Klues K, Xing G, Lu C. Link layer support for unified radio power management in wireless sensor networks. In *International Conference on Information Processing in Sensor Networks (IPSN)*, 2007.
35. Madden S, Franklin M, Hellerstein J, Hong W. The design of an acquisitional query processor for sensor networks. In *SIGMOD*, 2003.
36. Ramanathan N, Yarvis M, Chhabra J, Kushalnagar N, Krishnamurthy L, Estrin D. A stream-oriented power management protocol for low duty cycle sensor network applications. In *Proceedings of the Second IEEE Workshop on Embedded Networked Sensors (EmNetS-II)*, 2005.
37. Chipara O, Lu C, Roman G-C. Efficient power management based on application timing semantics for wireless sensor networks. In *International Conference on Distributed Computing Systems (ICDCS)*, 2000.
38. Klues K, Hackmann G, Chipara O, Lu C. A component-based architecture for power-efficient media access control in wireless sensor networks. In *SenSys*, 2007.

Authors' Biographies



Guoliang Xing is an Assistant Professor of Computer Science at City University of Hong Kong. He received the B.S. degree in Electrical Engineering and the M.S. degree in Computer Science from Xi'an Jiaotong University, China in 1998 and 2001, respectively. He received the Ph.D. degree in Computer Science from Washington University in St. Louis in 2006. He served on a number of technical program committees and held several workshop organization positions such as Program Co-chairs of the First ACM International Workshop on Integrated Heterogeneous Sensor and Actuator Networks (HeterSenet) 2008 and the Fifth Workshop on Wireless Ad hoc and Sensor Networks (WWASN) 2008. His research interests include wireless sensor networks, mobile computing, and networked embedded systems.



Mo Sha received his B.E. degree in Computer Science and Technology from BeiHang University, Beijing, China, in 2007. He is currently a Master of Philosophy candidate in Department of Computer Science at City University of Hong Kong, Hong Kong, China. His research interests include wireless sensor networks, ad hoc networks, and

active service.



Greg Hackmann received the B.S. degree in computer science from Washington University in St. Louis in 2004. He is currently pursuing a doctoral degree in computer science in the Computer Science and Engineering Department at Washington University. His current research interests include power management and middleware design for wireless

sensor networks.



Kevin Klues is a researcher in the Stanford Information Networks Group at Stanford University. He received his Masters degree from Washington University in St. Louis in May 2007 and his Bachelors degree from Rose-Hulman Institute of Technology in November 2002. From February 2003 to August 2005 he worked as a researcher in the Telecommunication Networks Group at the Technical University of Berlin, and before that spent time working at Motorola as well as Johns Hopkins Applied Physics Lab as an electrical engineer. His research interests primarily include techniques for achieving ultra-low-power operation in deeply embedded systems. He is a lead developer of the TinyOS project and responsible for the design and implementation of its resource and power management subsystems.



Octav Chipara received his B.S. in computer science from Indiana University Bloomington. He is currently a PhD candidate at Washington University in St. Louis. His research interests include real-time communication protocol for sensor networks, power management for embedded systems, and sensor networks for patient monitoring.



Chenyang Lu is an Assistant Professor of Computer Science and Engineering at Washington University in St. Louis. He received the B.S. degree from University of Science and Technology of China in 1995, the M.S. degree from Institute of Software, Chinese Academy of Sciences, in 1997, and the Ph.D. from University of Virginia in 2001, all in computer science. He received an NSF CAREER Award (2005) and a Best Paper Award from International Conference on Distributed Computing in Sensor Systems (DCOSS'06). He is author and co-author of more than 70 refereed publications. He also served on a number of technical program committees and held several conference organization positions such as Program Chair of RTAS'08, Sensor Networks Track Chair of RTSS'07, Demo Chair of SenSys'05, and Program Chair of WPDRTS'06 and FeBID'07. He is Associate Editor of ACM Transactions on Sensor Networks and International Journal of Sensor Networks, and Guest Editor of the Special Issue on Real-Time Wireless Sensor Networks of Real-Time Systems Journal. His research interests include real-time embedded systems and wireless sensor networks. He is a member of ACM and IEEE.