

Low-Power Downlink for the Internet of Things using IEEE 802.11-compliant Wake-Up Receivers

Johannes Blobel*, Vu Huy Tran[†], Archan Misra[†], Falko Dressler*

*School of Electrical Engineering and Computer Science, TU Berlin, Berlin, Germany

[†]Singapore Management University, Singapore

{blobel,dressler}@ccs-labs.org, {hvtran,archanm}@smu.edu.sg

Abstract—Ultra-low power communication is critical for supporting the next generation of battery-operated or energy harvesting battery-less Internet of Things (IoT) devices. Duty cycling protocols and wake-up receiver (WuRx) technologies, and their combinations, have been investigated as energy-efficient mechanisms to support selective, event-driven activation of devices. In this paper, we go one step further and show how WuRx can be used for an efficient and multi-purpose low power downlink (LPD) communication channel. We demonstrate how to (a) extend the wake-up signal to support low-power flexible and extensible unicast, multicast, and broadcast downlink communication and (b) utilize the WuRx-based LPD to also improve the energy efficiency of uplink data transfer. In addition, we show how the non-negligible energy overhead of conventional microcontroller based decoding of LPD communication can be substantially reduced by using the low-power universal asynchronous receiver/transmitter (LPUART) module of modern microcontrollers. Via experimental studies, involving both a functioning prototype and larger-scale simulations, we show that our proposed approach is compatible with conventional WLAN and offers a two-orders-of-magnitude improvement in uplink throughput and energy overheads over a competitive, IEEE 802.11 PSM-based baseline. This new LPD capability can also be used to improve the RF-based energy harvesting efficiency of battery-less IoT devices.

I. INTRODUCTION

With the introduction of low-power Internet of Things (IoT) concepts, the use of a wake-up receiver (WuRx) has become popular as an efficient way of mitigating the complexity of duty cycling-based communication protocols and reducing the power consumption for intermittent tasks (such as synchronization and reconfiguration) [1]–[3]. Following this WuRx concept, an IoT node is equipped with an additional ultra low-power (typically a few μW) radio receiver that remains active continuously with only negligible drainage of the battery, and is used only to wake up the main radio for event-driven communication [4], [5]. By dispensing with the need for periodic wake-ups, this paradigm overcomes the energy-vs.-latency trade-off inherent in duty cycling protocols. Moreover, event-driven wake-up radios have also been combined with duty cycling protocols to achieve even lower overheads for low bit-rate IoT communication [6], [7]. In multi-stage WuRx design, the wake-up signal is modulated using schemes such as on-off keying (OOK) or frequency shift keying (FSK) and includes a device-specific pattern or address to wake up only specific nodes [8]–[10].

In this paper, we propose a significant expansion in the use of WuRx capabilities, from merely activating the primary high-powered radio to enabling a more comprehensive low power

downlink (LPD) communication channel that can be used by a wireless LAN (WLAN) access point (AP) to interact with a set of nearby IoT devices. More specifically, we extend the WuRx mechanism for energy-efficient realization of two novel IoT control and communication capabilities:

Low-power downlink communication to IoT nodes: We shall show that, by generalizing the wake-up signal, the modulated data transmitted to a WuRx radio can be used not just to target a specific radio receiver, but also to support energy-efficient transmissions of intermittent application-layer data to the IoT node while keeping the node’s main wireless transceiver in deep sleep. Examples of such application-layer data include the sending of an acknowledgement (ACK) to confirm receipt of prior transmissions, commands to modify the configuration of on-board sensors and instructions to modify the duty cycle or periodicity of the main radio. Of course, as the data rate on the wake-up radio is usually quite low, the WuRx-based downlink transfer is likely to be useful only for relatively low volumes of data communication.

Efficient scheduling of uplink data transfers: Besides transferring configuration data or commands to the IoT device, the wake-up signal can also be used to perform scheduling of the uplink data transfers by the main radio. In general, when multiple IoT nodes are present, the uplink transmissions require additional contention resolution or channel access functions, which increase the overall transmission energy overhead. We develop an alternative centrally controlled channel access (CCCA) approach, where the WLAN AP, which also initiates the wake-up call, performs uplink channel access and scheduling (in time-slotted fashion) on behalf of one or more low-power IoT devices. To be fully 802.11 standard-compliant, the AP uses the *duration* field in the WLAN header to reserve the channel for the targeted nodes. The wake-up signal itself can either be sent by a second radio or by using cross technology communication (CTC) capability [11], [12].

To support the above two desired capabilities, we shall introduce two additional novelties in the specification of the wake-up signal protocol and the processing sub-system on the IoT node. First, to support transmissions of downlink data to groups of nodes, our wake-up protocol supports multicast and broadcast-based node activation, in addition to the usual unicast communication on the low power downlink. Our extended wake-up signal employs a parsimonious encoding scheme that not only permits selective wake-up of targeted IoT devices, but

also allows each such device to readily compute its designated, *contention-free* transmission slot on the uplink channel. Second, to further reduce the energy consumption on the IoT device, we shall make intelligent use of the low-power universal asynchronous receiver/transmitter (LPUART) module widely available in modern microcontrollers [13]. This module permits the decoding of the incoming LPD message while allowing the main microcontroller to remain in deep sleep, and thus provides an even more power-efficient mechanism to configure the operational parameters of the IoT node.

We shall demonstrate how our LPD communication technique can be implemented using a commodity hardware wake-up receiver, in tandem with the inter-processor communication (IPC) mechanism of ARM-based microcontrollers. We shall quantify the resulting benefits using both small-scale experimentation with a real hardware prototype and larger-scale studies using state-of-the-art simulation techniques. In addition, as an application example, we shall show how the novel capability of WuRx-triggered contention-free uplink transmission can be utilized to support significant improvements in the energy harvesting efficiency of a battery-less wearable prototype that harnesses power from beamformed WLAN transmissions.

Our key contributions can be summarized as follows:

- *WuRx-based Low Power Downlink*: We propose and demonstrate the feasibility of extending the functionality of a wake-up receiver to support an LPD communication channel that is 802.11 compatible. We develop the corresponding data protocol specification (Section III-D) that allows the LPD channel to be used for unicast, multicast, and broadcast communication, enabling ultra-low power transfer of short control messages from a WLAN AP to multiple IoT nodes.
- *Repurpose WuRx for Efficient & Reliable Uplink*: We also demonstrate how the wake-up receiver can be used to support highly energy-efficient uplink data transfer, across a varying number of concurrently active IoT nodes, while maintaining compatibility with conventional WLAN traffic. The proposed mechanism (Section III-A) utilizes the AP to perform reservation of the wireless medium on behalf of the low-power IoT nodes, followed by a time slot-based contention-free uplink transmission of data.
- *Ultra-low Power LPD Message Processing*: We further study and demonstrate the feasibility of processing LPD messages, which often require $O(10\text{ ms})$ reception time, via the use of a commonly available LPUART module that avoids waking up the microcontroller. Using a prototype real-world implementation of a wearable IoT device, we show (Section V-A) that wireless reconfiguration of our device can be achieved with an energy overhead of $22.3\ \mu\text{J}$, i.e., with 46% lower power consumption than achievable with microcontroller-based decoding.
- *Demonstration of Superiority over Competitive Baselines*: We utilize simulation studies to empirically establish the superiority of our proposed techniques. In particular, our proposed approach achieves an $\sim 87\% - 186\%$ increase in uplink system throughput and a $\sim 50\% - 98\%$ decrease in energy-per-bit transferred over alternative 802.11ba or

802.11(PSM) protocols. We also embed this mechanism in a WLAN-based energy harvesting wearable device and show that it can be used to optimize the detection of the node positions and, therefore, the beamforming for mobile nodes.

Overall, we believe that our work is the first to advocate for repurposing a commodity WuRx as part of an ultra-low power and extensible communication protocol (for both downlink and uplink traffic) for future IoT devices.

II. RELATED WORK

Wake-up Receiver: To overcome the inherent trade-off between energy efficiency and delay of duty cycling based protocols in the past years the concept of wake-up receivers has gained an increased interest in the research community [1]–[3], [7], [9], [10], [14]–[17]. In [16], the authors demonstrate how a system based on wake-up receivers can improve the performance in terms of delay and power consumption compared to a duty-cycling based protocol. In [7], a combination of a duty-cycling with wake-up receivers is used to create ultra-light sensors for tracking bats in the wild. A wake-up based system that includes addressing capabilities has been introduced in [2]. This system includes a routing protocol that leverages the selectivity introduced by the addressing to significantly increase the network lifetime. As the concept has proved to be very successful in sensor networks, it has also been adapted to other use cases. At the moment, an IEEE work group is designing a new wake-up mode (802.11ba) for the IEEE 802.11 standard [18]. The centrally controlled channel access (CCCA) proposed in this paper is building upon these concepts. In contrast to this standard, however, we do not use the compatibility mechanism to only guard the transmission of the OOK signal but also to reserve the channel for the following uplink phase.

Wireless Power Transfer: One of the key use cases for a wake-up receiver-based LPD explored in this paper involves the operation of battery-less IoT devices based on RF energy harvesting. The idea of wireless charging of embedded devices has been explored for both (a) very short distances ($< 5\text{ cm}$), for example using the near-field transfer based on the Qi [19] standard, and (b) longer distances to charge ultra-low power passive radio frequency identification (RFID) tags [20]. WLAN-based power transfer has been explored in the PoWiFi prototype [21], which introduced the notion of using additional WLAN “power packet” transmissions, simultaneously across multiple idle WLAN channels, purely to transfer power to an ultra-low power wearable device that operates with very low duty cycles. The WiWear system [22], extends the PoWiFi concept to utilize electronically beamformed transmissions of such WiFi power packets. This was first explored in [23] to achieve over a 100-fold increase in power harvested by a wrist-worn wearable device at distances of 1-3 meters. We should, however, note the following two relevant limitations of such systems: (a) it requires the periodic generation of *ping* packets from a device to the AP to optimally measure the node’s position and control the beamforming; (b) while the devices use a separate low-power non-WLAN transceiver to




Fig. 1. System Architecture: A central AP can communicate to both standard IEEE 802.11 and LPD IoT devices using a high-performance and a low-power radio, respectively. The AP can also use RF beamforming to support RF energy harvesting.

transfer its sensor data, the uplink transmission is performed in best effort fashion with *no contention resolution or access control* and suffers from high packet loss when the WLAN channel is experiencing moderate to high utilization. The system introduced in this paper can solve these shortcomings.

III. SYSTEM ARCHITECTURE

While wake-up receivers have been initially used in sensor networks, where all communicating nodes are energy constrained, we believe that it is useful to use them also in combination with a central access point (with unlimited energy) that acts as a gateway. Current technologies have an inherent trade-off between energy consumption and performance: modern WLAN radios provide very high data rates but are ill-suited for battery powered IoT devices while low-power technologies, such as IEEE 802.15.4 (ZigBee) or IEEE 802.11ah, only support very low data rates.

We propose a new system that uses an additional ultra-low power wake-up receiver on the mobile nodes to mitigate this trade-off. This additional radio is used to realize a low power downlink (LPD) from the AP to the mobile device. This asymmetric operation is attractive as IoT applications are often characterized by limited downlink communication and higher rates of uplink transfer of sensor data. While the main radio can utilize a standard high-performance technology like 802.11 WLAN, its energy consumption is still reduced by (a) executing most of the communication over the LPD, and (b) improving the efficiency of WLAN-compatible uplink data transfer. These features minimize the duration for which the main radio needs to be powered on. Of course, the wake-up radio can continue to be used as before, to wake up remote nodes.

The basic architecture of our system is shown in Fig. 1. The AP is the central entity that coordinates all communication. The radio of the AP can either send legacy WLAN signals or, using CTC, an OOK modulated LPD signal. The AP performs multiple tasks that include detection of new nodes, coordinating the channel access in a way that is compatible to WLAN, and data collection from mobile nodes. The AP may also feature an antenna array for beamforming to enable RF energy transmission. The low-power sensor nodes are equipped with two radios: A normal high-performance, high-energy radio that is turned off most of the time and a low data rate, low power LPD radio that is constantly in receive mode.

A. Centrally Controlled Channel Access (CCCA)

To ensure compatibility with existing WLAN networks and to simplify medium access for the mobile nodes, we propose a technique also used in the 802.11ba standard: before a frame exchange using the LPD, the AP performs regular channel sensing and backoff procedures to gain access to the wireless channel using the 802.11 distributed coordination function (DCF). It then sends out a normal WLAN header with the *duration* field set to a sufficiently large value. Other WLAN stations receiving this header will update their network allocation vector (NAV) accordingly and will not transmit during this time. This way, we can make sure that during a frame exchange initiated by an AP no other WLAN stations disturb the communication. This ensures compatibility with existing networks and is very beneficial in low power scenarios since devices can assume a free channel and can omit energy-expensive channel sensing.

B. Wake-up and Data transmission

For the downlink, we propose to use the CCCA as described before. In addition to the normal wake-up mode used in IoTs protocols and in the 802.11ba standard, our LPD uses the OOK modulated signal as a general purpose downlink. The frame format described in Section III-D supports different addressing schemes and can be used to send arbitrary data to the nodes. This data can be received in an energy efficient way and without waking up the node from deep sleep.

Given the energy constraints on individual nodes, their uplink transmissions should also be as energy efficient as possible. With the CCCA technique, we can realize a simplified channel access, because the AP already made sure that the channel is free for a specified time period after the LPD packet. As an example, we propose a slotted time division multiple access (TDMA) scheme where the AP announces a number of slots in the LPD packet. Each node then utilizes its own unique ID to select its own uplink transmission slot. The assignment of IDs can be done by the AP during the discovery of new nodes.

C. RF Power Delivery

To show how such a low-power communication protocol can be applied in a challenging IoT scenario with very little energy available at the nodes, we integrated it with a system to harvest RF energy emitted from the AP. To this end, the AP is equipped with an antenna array using beamforming to steer the




Fig. 2. Frame format for low-power downlink. The numbers denote the length of the fields in bits. Dashed fields are optional

emitted RF energy towards the IoT nodes (the antenna array is also used to estimate the angle of arrival (AoA) of received packets). The AP can use the LPD downlink to optimize the energy transmission, by periodic redirection of the RF beam, if nodes are mobile (see Section IV).

D. Frame Format

Fig. 2 shows the structure of our proposed frame format designed to support a flexible LPD functionality. Before each LPD frame, the AP sends a WLAN header as explained in Section III-A. Each frame starts with a preamble that signals the start of a frame. It is followed by two bits determining the *frame type*: 0 \rightarrow a wake-up frame, 1 \rightarrow a wake-up frame with additional data & 2 \rightarrow a data frame. The frame type is followed by a single bit *ACK field* that can be used to acknowledge the successful reception of an uplink packet previously transmitted by the main radio. This *splitting* of data and acknowledgment path promotes even greater energy efficiency for reliable communication by allowing the main radio to be powered down even before the arrival of the ACK.

The frame then includes fields to support efficient, variable-length addressing, which allows a single LPD frame to wake up one or many receiving IoT nodes. Each frame first includes a 2 bit *address type*: 0 \rightarrow broadcast frame, 1 \rightarrow unicast frame and 2 \rightarrow multicast frame. For a broadcast frame, no further addressing or mask information is needed; similarly, for a unicast frame, the *address type* is followed by a 8-bit *address*. Only multicast transmissions require a 16-bit field comprising an *address* and a *mask*, similar to the concept in [9]. Accordingly, the shortest LPD frame is a simple, 13 bit broadcast wake-up frame, containing the preamble, frame type, ACK and address type fields. For frames with additional data, there is an optional *length field* indicating the amount of trailing data. The LPD header is followed by an application-specific variable-length *data* field.

E. Use of the LPUART Module

While our LPD concept allows the IoT node to efficiently receive and transmit data, it does not directly address another challenge: the relatively high energy overhead of the microcontroller while it is reading and processing the data transmitted over the LPD channel. In effect, using the WuRx as an LPD helps minimizing the communication overhead, but does not address the computation overhead. For example, if the AP transmits a 64-bit packet to the IoT node using a 4 kbit/s transmission, the microcontroller must stay awake for 16 ms to read all the bits. Thus, the energy cost incurred by keeping

the microcontroller awake to receive downlink data may easily exceed the energy needed to transmit the response data.

Based on this observation, we propose to leverage the inter-processor communication (IPC) mechanism in ARM-based microcontrollers to support data reception via the LPD while keeping the relatively power hungry microcontroller in *sleep* mode. The IPC mechanism is implemented using a LPUART module, which stays active while the processor is in deep-sleep mode. The LPUART receives data and detects if there is a match between the received byte and its address. Using the LPUART, the system only wakes up to read LPUART registers when either an “address-matched” or “data-received” event occurs. This way, the microcontroller stays powered off during the entire LPD data reception process, and is activated only after reception of the entire frame.

IV. APPLICATION EXAMPLE

To illustrate the usage of our LPD concept in greater depth, we created a prototype application making use of the LPD communication principles. In the following, we describe the application-specific protocol aspects. The mobile devices are supposed to measure acceleration data and send this data to a central AP. Even though the main radio we used (nRF24L01+) already has a low power consumption (57 mW RX, 39 mW TX), it is still not possible to use duty cycling techniques or a carrier-sense multiple access (CSMA) scheme as the energy overheads are still too high. We show that the use of an AS3933 as a general purpose low power receiver, together with an application-specific protocol based on our LPD concept, can further reduce this overhead. The LPD is used for the following functions of the system: (1) Initial discovery of nodes, (2) adaptive ping to optimize beamforming, and (3) requesting sensor data. We therefore extend the frame format from Section III-D to include a custom payload and define the following communication protocol that is illustrated in Fig. 3.

Initial Node Discovery: Before sending energy to devices and requesting sensor data, the AP has to learn about the devices in range and about their position. This is achieved by sweeping the area with an RF power beam for a certain period to ensure all devices could harvest at least a little energy. The problem of node discovery is that initially the number of devices is unknown and a schedule for the uplink can therefore not be calculated by the AP. Our LPD concept allows to use the additional data field in the low power downlink to communicate scheduling data to the nodes, such as the number of available slots or a per-slot transmission probability (depending on the density of the network, similar to the ALOHA-like behavior of LoRA nodes [24]). Also the multicast capability of the frame format can be used to split up the discovery process into multiple stages if many nodes are in communication range. Such algorithms are highly scenario-dependent; however, some examples can be found in the realm of RFID tag discovery [25], [26]. We propose to use a simple slotted TDMA-based scheme, as it does not require any further data exchange. For simplicity, we assume that each node has a unique ID which determines the slot in which a node answers.




Fig. 3. Overview of the communication scheme. The AP first performs a node CCCA and then send broadcast/unicast packets via the LPD link to the nodes. The nodes then power on their main radio and directly transmit because the channel is blocked.




Fig. 4. Frame format for the discovery of new IoT nodes by the AP.

The discovery process starts with a broadcast LPD packet with the command field set to *Discover* (illustrated in Fig. 4). The nodes wake up and enable their main radio and send back a *HELLO* packet within its own slot, informing the AP about the ID of the node. The AP uses the channel state information (CSI) of the reception to determine the AoA. This information is later used to control the beam direction during power transmission phase. Within the *HELLO* packet, the device may inform about the sensors it supports.

Adaptive Pinging: In our scenario, the nodes have no battery and operate solely via RF harvesting. The efficiency of this approach highly depends on the correct estimation of the device position by the AP. If a node changes its position, the AP will continue to transmit the energy beam at the wrong angle, until it re-estimates the device’s position after receiving a new packet. We use adaptive pinging to continuously re-estimate the AoA. The protocol supports a *ping request* frame, which will wake up the node and instruct it to send a packet via the main radio to the AP. The AP can then adapt the rate of such requests to the mobility of each device and therefore optimize the beamforming and energy harvesting process. The uplink packet can also include state information like the charging level of the nodes that can be used to further adapt the energy transmission.

Requesting Sensor Data: Once the device has harvested enough energy, it can activate its sensors and record the sensed data. To receive this data, the AP sends a *data request* packet that includes one or more sensor IDs. A node transfers the requested data using the main radio.

V. EVALUATION

We now present results demonstrating various facets of LPD performance, using both a hardware prototype and extensive network simulations.

A. Hardware Measurements

We validate our proposed low-power downlink using a self-developed embedded system (shown in Fig. 5) including

four main components: (1) A main board equipped with an STM32L053 microcontroller which supports deep-sleep (or stop mode) with sub- μ A power consumption; (2) a wake-up receiver board which is tuned to work with 2.4 GHz band (same band as used for WLAN); (3) a main RF board (nRF24L01+) which is a commercially available RF transceiver board working at 2.4 GHz with a maximum bit rate of 2 Mbit/s; and (4) a sensor board which is currently equipped with a LIS3DH accelerometer. All the mentioned components are connected to the controller board via an SPI interface.

We implemented the AP using the WARP(v3) experimental board [27]. We use the open-source WARPLab as the experiment platform in Matlab environment; however, we integrate an FPGA-based packet detector to support real-time detection of nRF24L01+ packets. We use a data acquisition device (NI DAQ 6003) to capture the power consumption of the system during the lifecycle of both downlink and uplink communication. The device supports 100 kS/s with a resolution of 16 bit, which translates into a granularity of 14 mV. To measure the device’s power consumption, we measure the voltage drop across a high precision 10Ω resistor which is connected in series with the power pin of the device.

We conduct the experiment with two different implementations of LPD including: (1) a naive implementation where the device stays active, after it has been woken up by the AP, to read the pin level (presented on the DAT pin of the AS3933) to receive further downlink data transmitted by the AP. As the downlink communication via the wake-up receiver is slow (a maximum of 8 kbit/s if the Manchester coding is not used), the microcontroller needs to stay active for multiple milliseconds. (2) an LPUART-based implementation, where the LPUART is repurposed to receive data from the AP while the device is sleeping. While LPUART was introduced in ARM-based processors to support multi-processor communication, our approach allows the microcontroller to remain asleep instead of reading the downlink data bit by bit.

This is important because otherwise the energy for receiv-




Fig. 5. Our hardware prototype includes a main board, a wake-up receiver board, a main RF board, and a sensor board

ing via the LPD link can quickly become greater than the energy required for the main radio. Consider an STM32L053 microcontroller with an active power consumption of ≈ 5 mW at 2.5 V. The microcontroller spends ≈ 80 μ J to receive a 64-bit downlink packet. To compare it with the transmission of the response packet via the main radio, we assume the main wireless transceiver is a nRF24L01+ with a Tx power consumption of ≈ 28 mW (at 2.5 V). The longest possible packet (32 B) lasts only ≈ 0.3 ms (including transition time and overhead), which results in only ≈ 10 μ J. This relative comparison shows that the use of the microcontroller to poll the pin level to receive LPD data is inefficient.

We used the WARP board to transmit a wake-up pattern and two data bytes indicating a data request. To enable the data to be received by the LPUART, each byte is wrapped by a start bit and two stop bits (UART format). We intentionally add another stop bit to make sure it generates sufficient separation between two consecutive bytes. After the preamble (which is needed by the AS3933 to detect the data rate), 10 dummy bits of ones are padded to force the LPUART to IDLE state before the meaningful bits arrive.

Fig. 6 shows the device’s energy consumption during an LPD transaction. At 4 ms, the wake-up receiver starts to receive the incoming signal including the preamble. The actual data reception starts at 12 ms. When using the microcontroller to receive the incoming data, the power consumption rises to 3 mW during the complete reception time. In the LPUART mode, however, the microcontroller only has to be activated very briefly to configure the LPUART and to fetch the received bytes at times 12 ms and 16 ms (seen as the short spikes in Fig. 6). After the reception, the microcontroller activates the main radio and sends back a packet to the AP—this causes the high spike at 20 ms. The total energy consumed during this frame exchange is 41.5 μ J using the naive approach and only 22.3 μ J when using the LPUART approach.

B. Comparison with 802.11 Energy Saving and 802.11ba

To demonstrate the benefits of our proposed LPD protocol, we perform a simulation-based comparison of LPD vs. two other wireless communication protocols: Normal

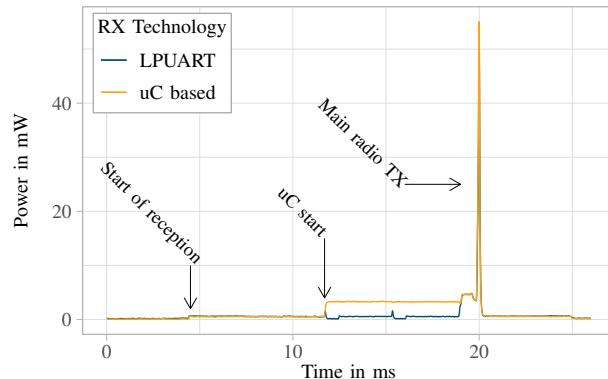


Fig. 6. Power consumption of prototype (with & without the LPUART) during reception of LPD frame and subsequent uplink transmission of a packet over the main radio.

802.11 WLAN with energy saving enabled (duty cycling) and 802.11ba. The LPD protocol is based on Section IV but without the adaptive pinging. To assess the performance of the different protocols, we analyze 1) the time it takes to request data from all nodes, 2) the fraction of successful requests, and 3) the energy efficiency of the protocols. The simulation model is based on the OMNeT++ network simulator version 6 and the INET wireless framework in version 4.2. The wireless channel was simulated using the `DimensionalAnalogModel`, which is well suited to analyze cross technology interference and coexistence.

For the WLAN-based scenario, the sensor nodes employ the 802.11 power save mode [28]: after a timeout where no data is received, a station (STA) goes to sleep mode and signals that to the AP. If the AP receives a new packet, it buffers the packet and indicates the availability of new data in the traffic indication map (TIM) of beacons sent out periodically (at 100 ms intervals). The nodes wake up to receive the beacons; if new data is available they request it from the AP. The 802.11ba WLAN standard draft [18] uses wake-up receivers to wake up STAs that have new data buffered at the AP. In this standard the wake-up signals are also prepended by a normal WLAN header to ensure compatibility with existing networks.

The simulation model in this section implements the following protocol:

- 1) In the start phase the AP learns about the surrounding nodes as explained in Section IV
- 2) The AP then requests sensor data from all nodes by sending a *Request* packet either
 - a) via broadcast to all nodes (The uplink channel access is done in a slotted TDMA fashion for the LPD protocol or using the 802.11 DCF for the WLAN based protocols), or
 - b) via unicast to each node consecutively.

Table I lists the most important simulation parameters. We compare the performance of the three protocols for different numbers of nodes in the network, different request intervals, and with unicast or broadcast request. In these simulations, we assume that there is no additional background WLAN traffic.




Fig. 7. Duration of the request phase (time to get sensor data from all nodes)




Fig. 8. Ratio of successful request phases to planned number of requests (depending on request interval); if one request phase takes longer than the request interval, not all requests can be fulfilled

Each configuration was simulated 10 times with a random placement of nodes around the AP.

Fig. 7 shows the mean duration of one request phase, i.e., the time it takes to collect the sensor data from all nodes. We can see that LPD has the lowest latency among the three protocols. When polling the nodes consecutively (as illustrated in Fig. 3) it takes 8 ms to query 4 nodes and 78 ms for 32 nodes. As expected this grows linearly (note the logarithmic scale), as each request takes the same amount of time. The wake-

TABLE I
SIMULATION PARAMETERS FOR THE PROTOCOL SCENARIO.

Description	Values	Unit
Simulation time	20	s
Number of sensor nodes	4, 8, 16, 32	nodes
Request Mode	Broadcast, Unicast	
Protocol	WLAN, 802.11ba, LPD	
Main radio	nRF24L01+, CC3200	
Request Interval	33, 100, 1000	ms
Slot length	1.2	ms

up based WLAN (802.11ba) shows slightly higher request duration (13–123 ms) because for each request the node has to be woken up, query the request packet from the AP and send back the sensor data. Even though the physical layer data rate for WLAN is much higher, this overhead and the additional waiting times caused by the DCF channel access scheme lead to a longer duration to query all nodes. For the native WLAN PSM mode, the duration is much higher. This is caused by the duty cycling approach, which requires the AP to wait until the transmission of the next beacon to notify STAs about the pending request packet. For unicast requests, this requires >100 ms to get the data from one node, resulting in large query times (378–3187 ms).

When using the broadcast approach, we can further reduce the time it takes to query all data. Here we only send one request packet to all nodes and receive the data using a TDMA approach from all nodes. The LPD protocol uses fixed slots (in our example 1.2 ms) for the uplink and does not require any additional channel sensing. Even though the WLAN-based approaches also transfer the data response packets to the network interface at the same time slots, the randomized nature of the channel access can cause collisions and require retransmissions when multiple nodes are present. For the native WLAN PSM protocol, the main determining factor of the request phase duration is the beacon interval. The time to transmit the data via the uplink is negligible compared to this. When using broadcast-based requests over the LPD, we were able to poll the data from all 32 nodes 435 times within 20 s—this represents > 186% improvement over the native WLAN PSM which completed such polling only 152 times, and 87% improvement over 802.11ba with 232 requests.

The duration of one request phase has an influence on how often the nodes can be polled. Fig. 8 shows the ratio of successful requests during the simulation time (20 s) and the number of planned requests (simulation time/request interval). Due to the low latency of the LPD protocol, it is able to satisfy all requests in almost all configurations, with the request satisfaction rate dropping from 100% only at the highest request rate (30 Hz) and for the cast of unicast retrievals or a larger number of IoT nodes. The slightly higher request duration of the 802.11ba approach leads to a reduced number of successful requests, especially when the number of IoT nodes is large. The WLAN approach shows the worst performance, with the satisfaction rate falling well below 50% in many cases.

Finally, we measured the energy consumed by the nodes over the simulation period. The values for the power consumption of the nodes during sleep, idle, transmission, and reception are taken from the data sheets of the chips used. To have a fair comparison, we used the same values for all three protocols. Fig. 9 shows the mean energy per transferred byte on the uplink channel. The energy consumed by a node is the energy required to send the data plus the energy required for communication overhead and idle time (the wake-up receiver is always on for the 802.11p and LPD scenarios). This is why the energy per byte is higher for the scenarios with a lower request rate (where less data is transmitted to the AP) even though the total energy

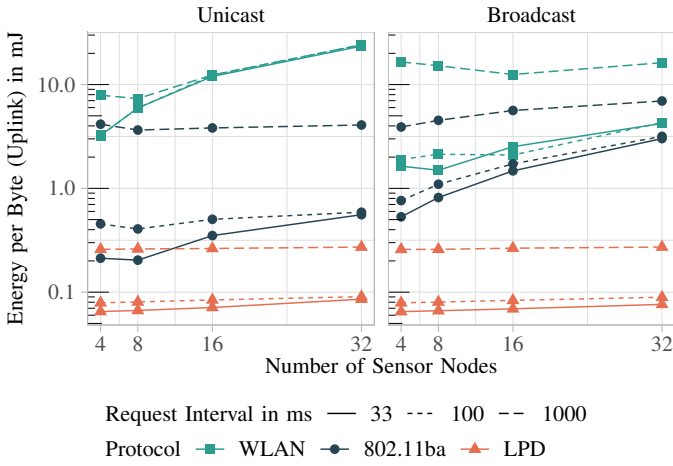


Fig. 9. Energy required to transmit one byte from the nodes to the AP; this includes overhead caused by the underlying communication protocol

(not shown here) is lower. The duty cycling approach of native WLAN PS shows the highest energy consumption (absolute and relative) due to the large overhead caused by the frequent wake-ups to receive the beacon from the AP. Additionally, the lower amount of total transferred data is lower for WLAN, thereby further increasing the increased energy consumption per byte.

The LPD protocol shows similar energy consumption for unicast and broadcast requests. This is due to the TDMA uplink scheduling, which allows nodes to power up their main radio *only* during their designated transmission slot, independent of the request mode or the number of other nodes. The 802.11ba approach, however, shows a significant increase in energy consumption when using broadcast requests. If all nodes are woken up at the same time, the DCF channel access causes many collisions, increasing the overall transfer duration and the resulting energy overhead. For 32 nodes and a request interval of 33 ms the LPD protocol can reach a power efficiency of 0.07 mJ/B which is only 1.6% of the power required by the WLAN protocol (4.22 mJ/B) and 2.3% of the power required by 802.11ba (3.01 mJ/B).

C. Coexistence with Other WLAN Devices

As described in Section III-A, the AP is responsible for reserving the channel for the low power nodes by first sending a WLAN header with the *duration* field set to an appropriate value. This ensures that our proposed LPD protocol is able to reduce the channel access overhead for the IoT nodes, while remaining compatible to existing WLAN devices. We now investigate the impact of our LPD protocol (and its low data rate) on existing or ongoing WLAN traffic.

We created a simulation scenario with one access point, 0–64 low power nodes and 2–3 normal WLAN devices (not connected to AP, but using Ad-Hoc/IBSS mode). The WLAN devices try to saturate the channel by sending sufficient user datagram protocol (UDP) packets to one of the nodes (receiver). At the same time and independent of the WLAN transmissions, the access point performs a node discovery and then requests

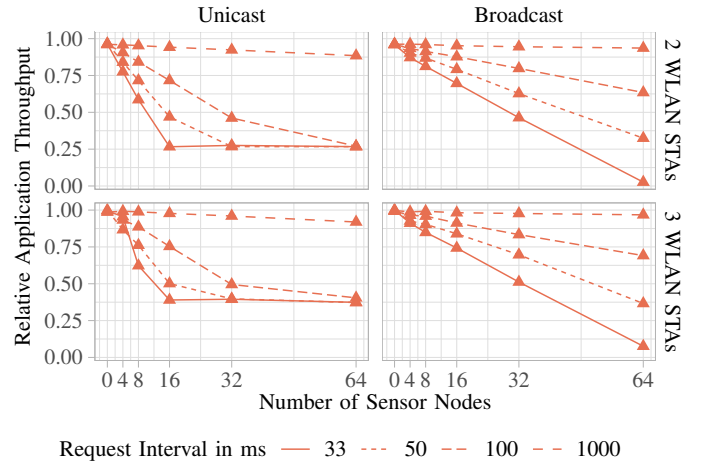


Fig. 10. Relative throughput of application: If no LPD devices are present the throughput is 1 (absolute value: 26.5 Mbit/s)

sensor data from the nodes at different intervals (33, 50, 100, 200, 1000 ms) as illustrated in Fig. 3.

Fig. 10 plots the relative throughput of the WLAN devices, where 1 is the throughput without LPD devices (26.5 Mbit/s). For all configurations, it can be seen that an increase in devices or an increase in the polling frequency (lower request interval) decreases the throughput of the ongoing/background WLAN traffic. This is anticipated due to the higher proportion of the shared wireless medium being reserved to enable the transmission of IoT traffic.

When comparing the results for unicast and broadcast data requests we can observe three effects. First, when the number of nodes is low (≤ 16), the broadcast approach does not reduce the WLAN throughput as much as the unicast approach. This is due to the fact that the broadcast request has less overhead as it only sends out the data request once instead of one request per node. Second, for a higher number of nodes, the unicast approach only limits the throughput down to 25% (1 WLAN sender) or 30% (2 WLAN sender). This difference to the broadcast approach is caused by the CCA scheme that requires the AP to first gain access to the channel using the 802.11 DCF. When using the unicast approach, for each node the AP first has to gain access to the channel, thus giving the other WLAN stations a chance to send as well. The differing share of the total channel capacity is due to the higher channel hold time (lower data rate) for LPD traffic and the well-known WLAN rate anomaly problem [29]. While a frame exchange of one 1500 B UDP packet takes only 274 μ s, one LPD slot reserves 500 μ s in this scenario.

This difference becomes even more visible when using the broadcast request where we can observe a third effect: When sending a broadcast request, the AP reserves the channel for an even longer duration = $slotLength * numberOfNodes$. For 64 nodes, LPD blocks the channel for 32 ms. Under very frequent polling of the sensor nodes (solid line), this causes the WLAN throughput to drop close to zero.




Fig. 11. Relative power received by the antenna array of the AP for different node positions. The Ping signals from the nodes are used to perform angle of arrival (AoA) measurements to update the beam direction when nodes move

D. Application of LPD in WiFi-based Energy Harvesting

Finally, we demonstrate that our LPD can be used to facilitate efficient wireless power transfer using WLAN. For this, we extended the hardware from Section V-A with a harvesting circuit that consists of a PCB antenna, a rectifier circuit and a voltage regulator to harvest the received energy and store it in a capacitor. The concept of using an antenna array to estimate the angle of arrival (AoA) and to use beamforming for an optimized transmission of energy using WLAN signals was already studied in [22]. The authors use a motion trigger to wake-up the microcontroller to transmit a “Ping” packet to an AP to re-calibrate the angle. This is a static mode that cannot adapt the rate of Pings to the motion of the sensor.

In this experiment, we show that LPD can provide an ultra-low power reactive Ping request method. We place the device at three different positions, which are right-most, center, and left-most, on a line slider in parallel with the antenna array of 4 elements. The slider is ≈ 90 cm far from the AP. We then change the position of the device from right to left. The device resides at each position for 30 s. The AP sends a Ping request every 2 s. Once the device is woken up by the request, it sends a dummy packet. The AP receives the uplink packet and estimates the AoA using the MUSIC algorithm [30] (Fig. 11 shows the pattern of the received signals as seen by the AP for the three different locations). The AP then changes the phase compensation for each antenna to update the beam direction.

Fig. 12 shows the harvested energy at each position and as well as during the move from one position to another. When the device moves to another position (or at the beginning of each run), the angle of the device is out-dated, and results in a significant drop in harvested energy (as seen at times 7 s, 40 s and 75 s). After the next Ping request, the angle is updated and the harvested energy is high again. Specifically at the beginning, the harvested energy rises from $4 \mu\text{W}$ to $40 \mu\text{W}$ after the AP receives a Ping response. The second transition shows even a higher rise ($20 \mu\text{W}$ to $150 \mu\text{W}$) in harvested energy thanks to the LPD enabled ping. This experiment shows that LPD enables smart WLAN-based energy harvesting even for mobile devices whose positions change over time.

VI. CONCLUSION

In this paper, we demonstrated how a wake-up receiver can be used to realize a general purpose low power downlink (LPD) communication channel. This allows an access point to

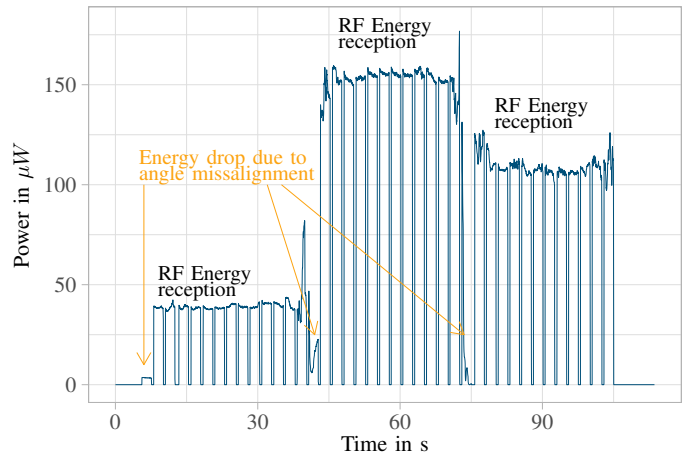


Fig. 12. The AP uses the LPD to request a sensor node to send a “Ping” packet at specific times, thereby enabling an updated estimation of the node’s AoA, thereby facilitating efficient beamformed power transfer

communicate with IoT devices in an energy efficient way and with high performance. In particular, our LPD concept allows ultra low-power downlink communication, e.g., to send configuration changes to (sleeping) Internet of Things (IoT) devices, as well as support contention-free channel access for the uplink communication, thereby avoiding the need for costly CSMA by each IoT devices. We introduced a flexible and efficient frame format that supports a wide range of possible use cases, ensures compatibility with existing wireless LAN (WLAN) networks and eases the channel access for the uplink. We developed a hardware prototype to show how such a technique can be implemented in an energy efficient way using the low-power universal asynchronous receiver/transmitter (LPUART) module of modern microcontrollers. We used this system in an example application, combining the low power communication with radio frequency (RF) energy harvesting capabilities, and studied the performance and the impact on existing networks, using both hardware experiments and simulation studies. Our results establish the superior features of the LPD concept and demonstrate how this new communication paradigm can overcome current limitations, enabling the deployment of ultra-low power (even battery-less) devices in key domains, such as industrial operations, while maintaining compatibility with existing technologies.

VII. ACKNOWLEDGMENTS

Research reported in this paper was supported in part by the German Federal Ministry of Education and Research (BMBF) in the context of the project *Energy efficient WLAN for IoT (EWI)* under grant number 01IS17046; as well as by the National Research Foundation, Singapore under its NRF Investigatorship grant (NRF-NRFI05-2019-0007). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

REFERENCES

- [1] R. Piyare, A. L. Murphy, C. Kiraly, P. Tosato, and D. Brunelli, "Ultra Low Power Wake-Up Radios: A Hardware and Networking Survey," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 4, pp. 2117–2157, Oct. 2017.
- [2] D. Spenza, M. Magno, S. Basagni, L. Benini, M. Paoli, and C. Petrioli, "Beyond Duty Cycling: Wake-up Radio with Selective Awakenings for Long-lived Wireless Sensing Systems," in *34th IEEE Conference on Computer Communications (INFOCOM 2015)*. Hong Kong, China: IEEE, Apr. 2015, pp. 522–530.
- [3] I. Demirkol, C. Ersoy, and E. Onur, "Wake-up receivers for wireless sensor networks: benefits and challenges," *IEEE Wireless Communications*, vol. 16, no. 4, pp. 88–96, Aug. 2009.
- [4] E. Alpman, A. Khairi, R. Dorrance, M. Park, V. S. Somayazulu, J. R. Foerster, A. Ravi, J. Paramesh, and S. Pellerano, "802.11g/n Compliant Fully Integrated Wake-Up Receiver With -72-dBm Sensitivity in 14-nm FinFET CMOS," *IEEE Journal of Solid-State Circuits*, vol. 53, no. 5, pp. 1411–1422, May 2018.
- [5] N. M. Pletcher, S. Gambini, and J. Rabaey, "A 2GHz 52 μ W Wake-Up Receiver with -72dBm Sensitivity Using Uncertain-IF Architecture," *IEEE Journal of Solid-State Circuits*, vol. 44, no. 1, pp. 269–280, Jan. 2009.
- [6] J. Blobel, C. Sommer, and F. Dressler, "Protocol Options for Low Power Sensor Network MAC using Wake-up Receivers with Duty Cycling," in *IEEE International Conference on Communications (ICC 2016)*. Kuala Lumpur, Malaysia: IEEE, May 2016, pp. 3925–3930.
- [7] F. Dressler, S. P. Ripperger, M. Hierold, T. Nowak, C. Eibel, B. Cassens, F. Mayer, K. Meyer-Wegener, and A. Koelplin, "From Radio Telemetry to Ultra-Low-Power Sensor Networks: Tracking Bats in the Wild," *IEEE Communications Magazine (COMMAG)*, vol. 54, no. 1, pp. 129–135, Jan. 2016.
- [8] C. Salazar, A. Cathelin, A. Kaiser, and J. Rabaey, "A 2.4 GHz Interferer-Resilient Wake-Up Receiver Using A Dual-IF Multi-Stage N-Path Architecture," *IEEE Journal of Solid-State Circuits*, vol. 51, no. 9, pp. 2091–2105, Sep. 2016.
- [9] J. Blobel and F. Dressler, "Sender-Triggered Selective Wake-Up Receiver for Low-Power Sensor Networks," in *36th IEEE Conference on Computer Communications (INFOCOM 2017), Demo Session*. Atlanta, GA: IEEE, May 2017.
- [10] J. Blobel, F. Menne, D. Yu, X. Cheng, and F. Dressler, "Low-power and Low-delay WLAN using Wake-up Receivers," *IEEE Transactions on Mobile Computing (TMC)*, 2020, online first, doi:10.1109/TMC.2020.3030313.
- [11] Z. Li and T. He, "WEBee: Physical-Layer Cross-Technology Communication via Emulation," in *23rd ACM International Conference on Mobile Computing and Networking (MobiCom 2017)*. Snowbird, UT: ACM, Oct. 2017, pp. 2–14.
- [12] M. Schulz, J. Link, F. Gringoli, and M. Hollick, "Shadow Wi-Fi: Teaching Smartphones to Transmit Raw Signals and to Extract Channel State Information to Implement Practical Covert Channels over Wi-Fi," in *16th ACM International Conference on Mobile Systems, Applications, and Services (MobiSys 2018)*. Munich, Germany: ACM, Jun. 2018.
- [13] *Ultra-low-power ARM Cortex-M0+ MCU with 64 Kbytes Flash, 32 MHz CPU, USB, LCD*, STMicroelectronics, Jun. 2019, rev. 9.
- [14] M. Bouraoui, I. Barra, and M. Masmoudi, "Study on wake-up receiver design for IoT applications," in *IEEE International Conference on Design & Test of Integrated Micro & Nano-Systems (DTS 2019)*. Gammarth-Tunis, Tunisia: IEEE, Apr. 2019.
- [15] M. C. Caballé, A. C. Augé, E. Lopez-Aguilera, E. Garcia-Villegas, I. Demirkol, and J. P. Aspas, "An Alternative to IEEE 802.11ba: Wake-Up Radio With Legacy IEEE 802.11 Transmitters," *IEEE Access*, vol. 7, pp. 48 068–48 086, Jan. 2019.
- [16] J. Oller, I. Demirkol, J. Casademont, J. Paradells, G. U. Gamm, and L. Reindl, "Has Time Come to Switch From Duty-Cycled MAC Protocols to Wake-Up Radio for Wireless Sensor Networks?" *IEEE/ACM Transactions on Networking (TON)*, vol. 24, no. 2, pp. 674–687, Apr. 2016.
- [17] S. Tang and S. Obana, "Tight Integration of Wake-Up Radio in Wireless LANs and the Impact of Wake-Up Latency," in *IEEE Global Communications Conference (GLOBECOM 2016)*, Washington, D.C., Dec. 2016, pp. 1–6.
- [18] IEEE, "P802.11ba/D2.0 - IEEE Draft Standard for Information Technology—Telecommunications and Information Exchange Between Systems Local and Metropolitan Area Networks—Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment: Wake-up radio operation," IEEE, Std 802.11ba-2019, Jan. 2019.
- [19] X. Liu, "Qi Standard Wireless Power Transfer Technology Development Toward Spatial Freedom," *IEEE Circuits and Systems Magazine*, vol. 15, no. 2, pp. 32–39, 2015.
- [20] D. J. Yeager, P. S. Powlledge, R. Prasad, D. Wetherall, and J. R. Smith, "Wirelessly-Charged UHF Tags for Sensor Data Collection," in *IEEE International Conference on RFID*. Kyoto, Japan: IEEE, Mar. 2008.
- [21] V. Talla, B. Kellogg, B. Ransford, S. Naderiparizi, S. Gollakota, and J. R. Smith, "Powering the next billion devices with wi-fi," in *11th ACM Conference on Emerging Networking Experiments and Technologies (CoNEXT 2015)*. Heidelberg, Germany: ACM, Dec. 2015.
- [22] V. H. Tran, A. Misra, J. Xiong, and R. K. Balan, "WiWear: Wearable Sensing via Directional WiFi Energy Harvesting," in *17th IEEE International Conference on Pervasive Computing and Communications (PerCom 2019)*. Kyoto, Japan: IEEE, Mar. 2019.
- [23] W. Huang, H. Chen, Y. Li, and B. Vucetic, "On the Performance of Multi-antenna Wireless-Powered Communications With Energy Beamforming," *IEEE Transactions on Vehicular Technology (TVT)*, vol. 65, no. 3, pp. 1801–1808, Mar. 2016.
- [24] L. Beltrame, A. Mahmood, P. Osterberg, and M. Gidlund, "LoRa beyond ALOHA: An Investigation of Alternative Random Access Protocols," *IEEE Transactions on Industrial Informatics*, 2020.
- [25] C. Law, K. Lee, and K.-Y. Siu, "Efficient memoryless protocol for tag identification (extended abstract)," in *4th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIALM 2000)*. Boston, MA: ACM, Aug. 2000.
- [26] J. Ryu, H. Lee, Y. Seok, T. T. Kwon, and Y. Choi, "A Hybrid Query Tree Protocol for Tag Collision Arbitration in RFID systems," in *IEEE International Conference on Communications (ICC 2007)*. Glasgow, Scotland: IEEE, Jun. 2007.
- [27] "Warp v3 kit," <https://mangocomm.com/products/kits/warp-v3-kit/>, accessed: 2020-08-14.
- [28] IEEE, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," IEEE, Std 802.11-2016, Dec. 2016.
- [29] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance Anomaly of 802.11b," in *22nd IEEE Conference on Computer Communications (INFOCOM 2003)*, vol. 2. San Francisco, CA: IEEE, Mar. 2003, pp. 836–843.
- [30] R. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Transactions on Antennas and Propagation*, vol. 34, no. 3, pp. 276–280, Mar. 1986.