

Network Interrupts: Supporting Delay Sensitive Applications in Low Power Wireless Control Networks

James Brown, Joe Finney, Christos Efstratiou, Ben Green,
Nigel Davies, Mark Lowton and Gerd Kortuem

Computing Department
Lancaster University
Lancaster, UK

{jb, joe, efstrati, greenb, nigel, lowton, kortuem} @comp.lancs.ac.uk

ABSTRACT

The importance in maintaining energy efficient communications in low power networks such as sensor and actuator networks is well understood. However, in recent years, a growing number of delay sensitive and interactive applications have been discovered for such networks, that are no longer purely limited to the data gathering model of sensor networks. Providing support application requiring low latency interaction in such environments without negatively affecting energy efficiency remains a challenging problem. This paper outlines the importance of this emerging class of application, discusses problems involved in supporting them in energy challenged environments, proposes a combined hardware and software mechanism based on heterogeneous wireless networking which works toward solving this problem, and goes on to evaluate this mechanism through experimental analysis. The paper concludes with a discussion of the applicability of the mechanism to typical application scenarios.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless communication; C.2.2 [Computer Systems Organization]: Computer-Communication Networks— Network Protocols

General Terms

Measurement, Experimentation.

Keywords

low power networking, heterogeneous devices, interactive applications, delay sensitive applications

1. INTRODUCTION

The NEMO project is exploring the use of ubiquitous technologies and embedded wireless systems in industrial workplaces. Industries such as construction and road maintenance are characterized by environments where manual labor is

performed in harsh environments with minimal or no IT support. Our aim is to explore the potential to improve the work practices, management and coordination of activities performed in such industries, by employing embedded wireless technologies. Indeed, we envisage a world where physical work artifacts such as tools, vehicles and workers are augmented with embedded cooperating mobile nodes featuring both sensors and actuators. These nodes can form ad-hoc networks, utilize their sensing capabilities in order to observe the activities performed, and collaboratively offer assistance to the worker crew, when necessary.

One of the primary areas of interest is the development of systems to support health and safety (H&S) in construction and road maintenance sites. Augmented artifacts offer the potential for automated assessment of compliance with H&S regulations and in-situ notifications of the workers when violations occur. The design of such systems is complicated by both the functional requirements of the targeted applications as well as the operational requirements imposed by the industry. Low latency communication is a principal requirement for the implementation of H&S applications over ad-hoc networks. Equally, the adoption of any such system by the industry requires long battery life with minimal maintenance. In this paper we present a communication mechanism that satisfies both these requirements. The mechanism relies on heterogeneous wireless networking, utilizing a combination of ultra low power radios and traditional short range radios. The main contribution of the work lies on our real world implementation of the system, and evaluation and analysis of its performance.

1.1 Motivation

In collaboration with a major construction company we have been investigating the use of embedded wireless systems to support their field workers. After a series of interviews and field studies, we have identified a number of scenarios where the use of embedded wireless systems would offer significant benefits to both the workers and the company.

1.1.1 Monitoring Hand-Arm Vibrations

One of the scenarios we have been investigating involves the development of a system to monitor workers' exposure to hand-arm vibrations (HAV) caused by vibrating tools such as drills and jackhammers. Long-term exposure to HAV can lead to serious health conditions such as "vibration white finger". In order to monitor the workers' exposure to HAV and avoid violation of related H&S regulations, we designed, implemented and field

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

CHANTS'07, September 14, 2007, Montréal, Québec, Canada.
Copyright 2007 ACM 978-1-59593-737-7/07/0009...\$5.00.

tested a prototype system for monitoring HAV exposure in road maintenance sites [1].

The HAV monitoring system comprises of a collection of mobile nodes that collaborate in order to observe the usage of vibrating machinery. Tools augmented with wireless sensor nodes collaborate with wearable personal NEMO devices in an ad-hoc manner, in order to record each worker's exposure to HAV. Communication between the nodes is mainly driven by the workers' activities: when workers operate vibrating tools (e.g. drills), the tool's sensor node delivers HAV exposure events to the worker's personal device, where they are being recorded and assessed against the H&S regulations. If a worker exceeds the amount of time prescribed by the regulations, an alert is raised and a notification is presented on the screen of their personal device.

1.1.2 Real-time Emergencies

In the construction industry there are numerous situations where instant alerts are required in order to avoid accidents. Common scenarios that can lead to serious injuries or death, include vehicles backing-up, where limited visibility could lead to accidents. In the road maintenance industry in particular, one of the important causes of serious accidents is the violation of traffic management signs/indicators by by-passing cars. Augmented vehicles and equipment (e.g. signs, traffic cones) offer the potential to improve the chances of avoiding an accident in these situations. Embedded wireless sensor nodes can collaboratively assess the conditions in a particular site and raise alerts in real-time (e.g. when a traffic cone is run over by a car). Such system would require: (i) timely detection of a critical situation or a violation of safety regulations, (ii) immediate alert of the involved personnel, (iii), and immediate actions to reduce the risks or effects of an incident, when possible.

1.1.3 Asset Management

In any industry, detailed record keeping of a company's assets is an important part of their administration. In the road maintenance industry, high mobility of assets (e.g. large number of mobile maintenance crews), along with large numbers of assets that require regular maintenance (e.g. street lights, signs), makes asset management especially challenging. The employment of embedded wireless systems for asset management would comprise of assets augmented with wireless sensor nodes in order to provide location tracking and detailed usage and maintenance history of each asset. The communication patterns expected by such systems involve long periods of inactivity, as such smart artifacts (tools, equipment, etc.) often spend large periods of time in storage, and are not required to interact in these periods.

1.2 Requirements

As illustrated in the previous sections the wireless sensor nodes that are embedded into vehicles and tools are required to support the diverse communication requirements of varying application scenarios. Communication patterns can vary significantly, from very long times of inactivity, as indicated by the asset management scenario, to real-time/short-delay responses to support emergency alerts. In tool usage scenarios, such as the HAV monitoring system, communication is mainly driven by human activity, and therefore communication patterns are hard to predict.

In addition to these diverse communication patterns, the realization of such system is further complicated by a number of operational requirements imposed by the nature of the construction/road maintenance industry. Specifically, maintenance of the wireless nodes should be minimal. Considering that a company would have to deploy hundreds or thousands of nodes, any regular maintenance, including recharging of wireless nodes, would impose an unacceptable overhead. Indeed, the construction company we collaborate with would favor low cost disposable nodes that have a battery life of at least a year.

Summarizing these findings, the key communication requirements for these systems are:

- R1.** Embedded nodes should maintain a long battery life (>1 year).
- R2.** Nodes should be able to operate without any supporting communication infrastructure.
- R3.** User-driven communication should maintain short delay responses (<1 sec).
- R4.** Emergency communication should be performed with minimal delay (<100ms).
- R5.** System should be optimized for the high variation of communication frequency: long periods of inactivity and short, unpredictable periods of low latency interaction.

2. NETWORK INTERRUPTS

2.1 The need for a different approach

The combined requirement of supporting both low power operation and short response times is highly challenging. It is widely known that the greatest energy drain for such devices is the operation of their wireless network interfaces. For example, a commonly used implementation of the IEEE 802.15.4 standard, the Chipcon CC2420 radio [2], represents the current state of the art in low power communications, and its power consumption remains over two orders of magnitude higher than that needed to support these applications for one year in the field in an always on configuration.

As such, common approaches to improving battery lifetime focus upon placing network interfaces and processors in sleep modes for a large proportion of their time, and awaking periodically to service application and communication needs. It is not at all uncommon for a deployed WSN device to remain dormant for several seconds at a time, in a 1% duty cycle (i.e. spending up to 99% of its time in a sleep mode). There are a number of issues with the use of such periodic sleep scheduling in wireless networks that are required to support applications as varied and interactive as those described earlier:

Firstly, although sleep scheduling dramatically improves energy efficiency, it does so at the cost of significant additional communication latency. Devices operating in this fashion cannot guarantee a response time better than their sleep interval. Furthermore, the latency patterns caused by sleep schedules can have undesirable effect on jitter. In the case of the interactive or emergency scenarios described previously, it is clear that such delays between nodes can be extremely prohibitive.

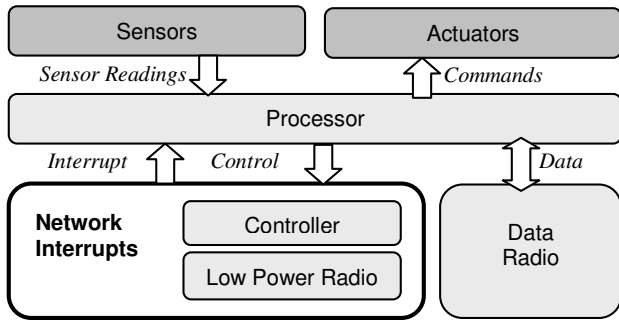


Figure 2: System Architecture

Secondly, the unpredictable traffic patterns characterized by short periods of interactivity and long periods of total inactivity (i.e. during storage) make the definition of a sleep schedule problematic. Simply enforcing shorter sleep intervals to meet low latency requirements would significantly reduce energy efficiency over the lifetime of the device, likewise defining longer intervals results in the inability to respond quickly to events and suffer from the problems of clock skew between devices, thwarting the potential energy saving benefits.

Finally, the highly mobile nature of these artifacts leads to the need for low overhead device and resource discovery. It is very common for devices in such networks to form and leave networks very frequently, as tools, workers and pieces of equipment are moved around the workplace. Discovering neighbors and performing the clock synchronization required to define sleep schedules in an energy efficient way is a significant problem. Without external timing information, a node joining a network would expect to remain active for at least half of the sleep period just to discover the presence of a neighbor. In the case of the 1% duty cycle example earlier, a node would consume *at least 100 times* the energy in that discovery period than in its normal operation. This cost can become dominant in such dynamic networks.

In summary, current schemes to enable long life in energy constrained environment come at the cost of an unacceptable degree of latency and loss in dynamicity. The only real solution to supporting delay sensitive applications is to reduce the duty cycle of the devices to meet application requirements, resulting in much wasted energy caused by nodes listening to an empty channel, and thereby reducing energy efficiency beyond what is practical with today’s battery technologies.

2.2 Ultra low power radios

Modern digital radios are highly sophisticated, with high bandwidth, strong channel coding, error detection, error correction and MAC mechanisms. Although highly beneficial in many situations, it is often overlooked that these features are the cause of much of the radio’s power consumption. Earlier, relatively simplistic radio designs exhibit significantly reduced power characteristics, particularly at the radio receiver. Commercial implementations of such radios are now becoming commonplace. One commercial example of this is the AM-HRR18 433Mhz receiver [3] (Figure 1). This is a modern implementation of the super-regenerative radio receiver, first patented in 1922. This radio costs only a few dollars, and operates at a power consumption of 70uA – *almost three orders of*




Figure 1: AM-HRR18 Super Regenerative Receiver

magnitude lower than the CC2420 IEEE 802.15.4 radio, which consumes approximately 19mA.

The AM-HRR18 radio in comparison to the CC2420 is very simple; no encoding and error correction is performed by the radio, leaving such tasks to be handled by an external processor. Furthermore, the significantly lower bandwidth (a meager 4kbps in comparison to 250kbps) and the huge power-on time of this receiver (Table 1) makes it inappropriate for systems requiring either regular sleeping cycles or high bandwidth. However, due to the extremely low power consumption of such radios and the fast transmitter power-up, they can make excellent “always-on”, secondary control channels for more powerful data radios, such as the CC2420 – to form a *heterogeneous MAC protocol*.

2.3 The NEMO hybrid radio device

We have designed and developed a prototype interactive sensor device for experimental research into supporting the NEMO scenarios described earlier. A key feature of this device is its split level design, based upon a hybrid high and low power processor and radio pairs arranged in a hierarchical fashion. The primary sensor system is built around a 32-bit NXP ARM7 LPC2136 [4] microcontroller and a Chipcon CC2420 802.15.4 radio. This is augmented with a low power Microchip 16f737 8-bit microcontroller and a low power AM-HRR18 433Mhz radio receiver module, coupled with a matched AM-RT5 radio transmitter (Figure 2).

The low power radio and microcontroller are always powered on in this design, and are used as a control channel for the primary radio and processor, which it can wake from sleep state via an external interrupt. Beacons can be sent at any time over the low power radio from any node to wake up the primary radio of another node in range, such that nodes only enter an active listening state when there is data to be received.

2.4 Beacon Transmissions

Care must be taken in how the beacon is transmitted. As most common RF channels are within the ISM band, heavy noise must be expected. However, there is a trade off here, since the more complex and error resilient a beacon is, the more complex the

Table 1: Key Characteristics of Hybrid Radios

	Bitrate	RX current	TX current	Power on delay
CC2420	250 kbps	18.8mA	17.4mA	1.3ms
AM-HRR18	4 kbps	70uA		2000ms
AM-RT5	4 kbps		4mA	10us

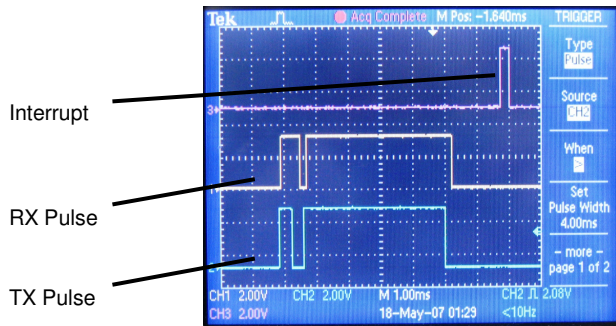


Figure 3: Beacon Modulation

process of detecting that beacon becomes, thereby adding to the energy cost of transmitter and receiver processing.

In our scheme we do not employ any framing, error correction or detection mechanisms for the beacons. A beacon is encoded purely as a single pulse. This minimizes the transmission and reception complexity, processing cost and transmission delay. Our early experiments have shown a short sync pulse followed by a HI pulse length of 5ms surrounded by LO pulses of 250uS is sufficient to separate a beacon from typical background noise, as illustrated in Figure 3.

The use of such a simple scheme negates the need for DSP processing on the receiver, and can be detected by an 8 bit microcontroller operating at only 32kHz and consuming 30uA – approximately the same processing power as a digital watch. Upon detection of such a pulse, the microcontroller wakes up the main processor from its sleep state via an interrupt, just in time to receive the incoming transmission on its primary radio.

2.4.1 Targeted Beacons

We envision scenarios where tens or even hundreds of NEMO devices will ultimately be within each other’s communication range. It is therefore important to ensure that beacons do not wake up all nodes within transmission range of the transmitter, but to target specific nodes. We support targeted beaconing with minimal additional complexity by utilizing pulse width modulation (PWM). Transmitters take a six bit hash of the target node’s MAC address and calculate a pulse-width of:

$$5 + \text{hash_fn}(\text{MAC address}) * 0.25 \text{ ms. Eq. 1}$$

This calculation yields a total wakeup delay of between 5ms and 20ms, can be easily decoded by the receiver, and can significantly reduce the number of devices unnecessarily awoken. Moreover, this PWM based scheme allows for a degree of error tolerance without the need for error control coding – e.g. a node could wakeup to a pulse which is ‘close’ to its hash (within a degree of tolerance), thus easing effects of radio noise. We are also investigating more robust alternatives to hashing based on local cluster based addressing.

2.4.2 Collision avoidance

For the low bitrate applications this work aims to support, we believe contention between nodes for the control channel is likely to be low, but uncontrolled access to the channel would make collisions inevitable. We utilize the always on nature of the receiver to employ a simple CSMA/CA strategy. Nodes will not initiate beacon pulses if they detect ongoing activity. We do not attempt to address issues of hidden transmitter problems, and

instead rely on higher level retransmissions to resolve these events.

2.4.3 Discussion

There are a number of benefits associated with this approach with respect to the issues described above. Firstly, there is the potential for power saving rising from the fact that receivers no longer need to periodically awaken to determine if there is data to receive. Secondly, the spontaneous nature of the beacons facilitates lower latency communications, since the low power receiver on the node would always be active. Thirdly, this approach does not suffer from the problems of clock skew, as no shared clock is needed between devices. Finally, the beaconing mechanism can be used to pre-empt service discovery phases, thus enabling low latency discovery with low energy overhead.

There are also a number of potential drawbacks to this design, the primary being the additional energy consumption, financial costs and hardware footprint associated with the secondary radio. We believe that the additional financial and spatial costs are minimal, and would be negligible if such radios were in wide spread use. To determine the energy costs and benefits, we go on to examine the accuracy and efficiency of the scheme in the next section.

3. EVALUATION

In this section, we evaluate the efficiency of using a secondary lower powered radio to provide network level interrupts for long lived control applications. We compare and contrast this system with periodic sleep schedules, the more traditional approach to achieving longevity within energy constrained environments, and investigate the relative performance of these schemes in terms of energy efficiency of receivers and transmitters, packet latency, and effects of background noise and communication range.

We choose to evaluate this approach through the use of laboratory experimentation. We believe that evaluating the system on real hardware gives us more accurate insights into how the system would perform in the field relative to an analytical or simulated study.

3.1 Overview of Testbed

The testbed used to evaluate the two approaches comprises of two prototype NEMO sensor devices (figure 4), one used for transmitting and the other for receiving. These nodes were both instrumented with ammeters capable of real-time logging of the devices’ power consumption.




Figure 4: Nemo Device

We decided to isolate the main processor from any monitoring instrumentation during each of our tests. This allowed us to only measure the cost associated with the communications. As a result a more accurate set of results are obtained as the costs associated with taking, recording and conveying measurements to the data logging system are excluded from our results. Furthermore, the selection of main processor is quite diverse from system to system, so only measuring the cost of communications, we believe, makes the results more directly comparable to other systems.

A PC is used for the purpose of data logging and is connected to each node. The data log records contain detailed timing and status information such as wakeup and sleeping times, received and missed packets. With this information, it was then possible to calculate offline the following.

- **Average Response Time.** The time taken between a data packet to be generated on the transmitting node and being received on the receiving node.
- **False Negatives.** The occasions where the transmitter was active, while the receiver was sleeping, indicating a missed wakeup.
- **False Positives.** The occasions where the receiver awakens, but no packets were received, indicating an erroneous wake-up.

In addition to the above statistics the energy consumption of both the receiver and transmitter is also recorded. To acquire this, two Fluke 189 data logging multi-meters were used, which record the current consumption averaged over 100ms, every 100ms. All measurement recorded were time stamped using a high frequency timer to enable event correlation.

All tests were carried out in an indoor office environment. Two nodes were placed on a lab bench at a distance of approximately one meter for energy consumption and latency testing. Range tests were conducted at a similar height but the distance was slowly increased to a maximum of thirty meters. The wake up radio used a ¼ length wave whip antenna vertically orientated.

3.2 Energy Utilization

To provide a fair evaluation of the network interrupt system, we compare it against a static sleep scheduling approach with a variety of duty cycles and data rates. This acts as a control for our experiment.

To perform the tests on a static scheduling system we configured the devices to sleep for a period of time then wake and communicate. The transmitter was configured to randomly generate packets at a configurable rate. All experiments used a fixed packet size of 60 bytes. Each packet was buffered until the next wake up period at which point all the packets currently stored were transmitted. The receiver was configured to wake up at each scheduled point and listen for 20ms. If a packet arrived during this time the 20ms counter was reset and the receiver continued to listen. Once the counter expired, the receiver would return to sleep until the start of the next wake up period.

It is worthy to note at this point that real world schemes based upon static scheduling approaches often implement a number of optimizations to improve energy efficiency. The radio sampling window is determined by a number of factors including radio

wake-up time, sampling time and system clock drift. All these factors can be optimized for specific architectures, as proposed by other works in the literature. Here we choose a fixed wake up window of 20ms to provide a common baseline for comparison to other systems, allowing us to compare directly to other schemes based on their ‘effective’ duty cycle.

To test the network interrupt mode, the transmitter was similarly configured to generate packets at a specified rate at random intervals. On the generation of a packet, the transmitter would signal the low powered radio to transmit a wake up pulse. The transmitter would then wait for a sufficient period of time for the pulse to propagate before enabling the main transmitter and sending the packet. The transmitting node repeated this for the duration of the test. The receiver system was configured to sleep until interrupted by the low power radio. On interrupt the receiver enabled its receiver and similarly waited 20 ms for the packet to arrive. After either a packet had arrived or the 20ms timer had lapsed the receiver returned to sleep mode awaiting the next pulse.

Each test was configured to last 300 seconds, a sufficient amount of time to produce stable results. Both nodes were powered from the same source to provide a controlled method of clock synchronization. Nodes were reset at the start of each test simultaneously to remove errors caused by clock drift. The sleep periods used during testing on the static scheduling system were set at 0.5, 1, 2, 4, 8, 16 and 30 seconds. We believe this covers the common configurations for most WSN applications. A sleep period longer than 16-30 seconds would seriously hamper interactivity whereas one of less than 0.5 seconds would impair battery life.

We tested these configurations with a range of network loads, ranging from 1200, 1000, 800, 600, 400, 300, 200, 100, 60, 30, 15, 10 and 6 packets per minute. This produced a packet every 50ms in our heavily loaded scenario and conversely one every 10 seconds in our lightly loaded scenario. These figures are high for our application domain but we believe these would lead to the most insightful results. Our nodes used a Poisson distribution for the inter-arrival times of each packet to better emulate the interactive traffic we expect in NEMO scenarios. All 104 tests were averages over three runs to improve their accuracy.

3.2.1 Receiver Energy Consumption Results

Our first set of results examines the power consumed by the receiver system. Figure 5 summarizes these results, illustrating the

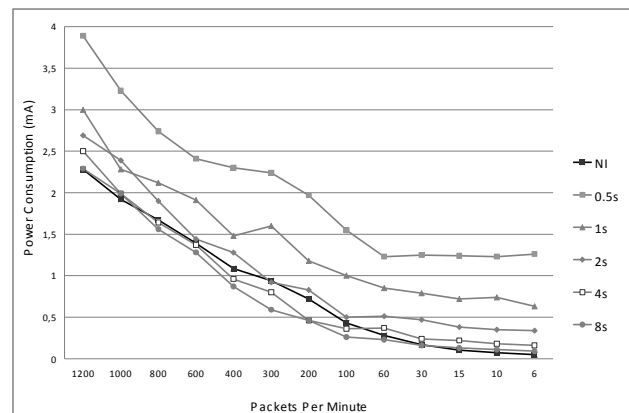


Figure 5: Energy Consumption at Receiver

relationship between power consumption (y axis) and network load (x axis) for the Network Interrupt scheme, and static scheduling schemes at sleep intervals from 0.5s to 8s.

We can see the results for the static scheduling mode are stacked based on wake up frequency. Although the actual power consumed for receiving N packets in any interval should remain constant regardless of duty cycle, other factors affect total consumption. The main influencing factor is the increase in receiver idling time as wakeup frequency increases, i.e. when the receiver wakes periodically, but there are no packets to be delivered.

The results for the network interrupt scheme show a different profile to that of static scheduling, with an rx energy consumption comparable to a static scheduling scheme with a sleep period of four seconds at loads of 100 packets per second. At reduced loads the schemes performance improves, exhibiting energy consumption comparable or better to an eight second duty cycle at 30 packets per second or less.

3.2.2 Transmitter Energy Consumption Results

In terms of transmission power, figure 6 summarizes the result. We can see the results for the static scheduling systems mirrors that of the receiver - stacked by wake up frequency due to the additional power consumed in initializing the radio more often at higher frequencies.

The results for the network interrupt scheme follow the same trend as the static scheduling system, but at a slightly higher consumption. This is due to the use of the additional radio, and the power consumed during the initialization of the main radio each packet, rather than once at the end of each duty cycle. We note that the static scheduling results are stacked by wake up frequency, but also contain some outlying points. These we believe can be accounted to the Poisson random number generator, generating slightly different distributions of packets across each of our tests. The results show that it is approximately 1mA more expensive to use the wake up system at higher loads. However at the more realistic loads of less than 100 packets per minute the cost is approximately 100-200uA greater than that of the static system.

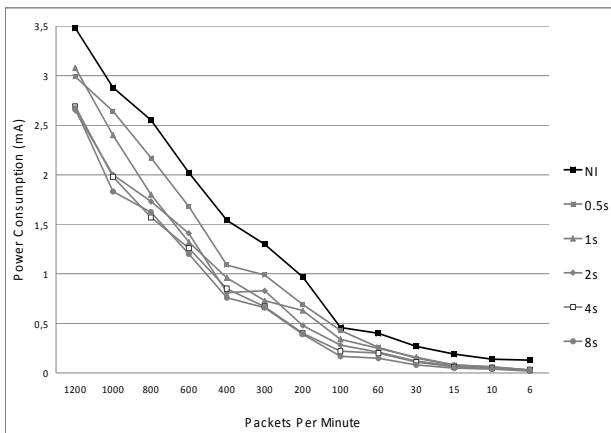


Figure 6: Energy Consumption at Transmitter

3.2.3 Packet Latency Results

Further results examine the scheme's response time, which we define as the time it takes for a packet once generated to be transmitted from source to destination. For the static system we expect the average response time to be half of the sleep schedule (the Expectation). The response time of the network interrupt system should be constant as this is the pulse time, which is constant, plus the time to send the packet which should also be constant as channel will be clear during our tests. Figure 7, presents the response time of the wake up system in comparison to half, one and two second sleep cycles. As we expected the response time of the static system is half the sleep period, however this does rise a little as network load drops. We attribute this to the fact that nodes spend more time awake under heavy network load, so packets generated at the start of a sleep window are only queued momentarily before transmission as the node has yet to sleep. The results for the network interrupt scheme showed an inverse reaction to load in comparison to static schedules. At moderate load levels of 400 and less the response time is constant at 15ms, however this rises to 19ms as our highest tested load of 2400 packets per minute. This is caused by packets being generated and queued whilst previous packets are still being sent.

These results have shown that the power consumption is comparable to duty cycling as the power consumed by the additional radio is offset by spending less time idling at the lower loads we would expect to see in the real conditions. However, the response time offered is many orders of magnitude lower than duty cycling.

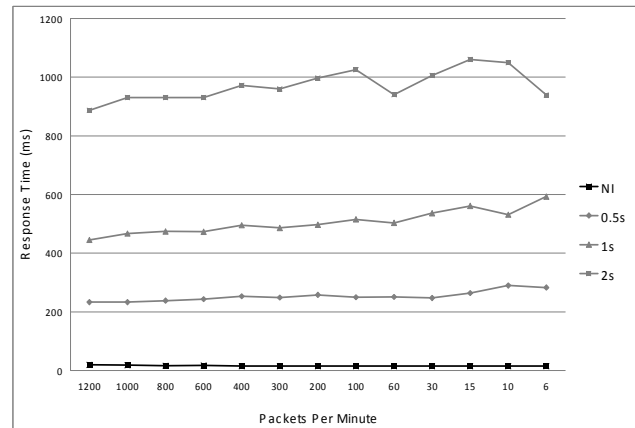


Figure 7: Response time of static schedules and network interrupt scheme at varying network loads.

3.2.4 Detection Accuracy and Noise

It would not be unreasonable to assume the lack of error control in our beacon mechanism could lead to a number of false positives and false negatives in the detection of beacons. During all our tests, we used timestamp comparisons to detect and record such erroneous events. Furthermore, to make our tests more realistic, we made no effort to clear the 433 MHz ISM channel (which was found to contain traffic from other users). In addition, we adopted a pessimistic view of the results, where all false negatives were attributed to the wake up radio and not from the data radio. We noted that over all the undertaken tests (24125 transmitted packets) we observed a total of 8 false positives and 7 false negatives. We feel this is a relatively small proportion of packets,

for all the potential benefits. Neither a false positive nor negative would typically cause a major failure. False positives result only in unnecessary wake up of a device wasting energy, whereas a false negative would result in higher layer retransmissions and therefore increased latency. No systematic patterns were discovered in these results.

3.2.5 Transmission Range Tests

A number of experiments were also undertaken to determine the effect of transmission range on beacon reception accuracy. We performed six trials at ranges from 5m to 30m. Each test ran for one minute and contained 600 beacons during each test. During these tests, the number of accurately detected wake events on the receiver was recorded. We found that no beacons were lost up to this range in any of our tests, nor were any false positives noted, despite the low power nature of the transmitter and receiver.

One further aspect of concern is matching the transmission range of the two radios. However, our experiences bring us to believe that such concerns could be engineered out. Firstly, modern radios (such as the 802.15.4 radios used in these experiments) often permit the setting of transmission output power in software, and as such can be configured to reduce its transmission range to match that of the wake-up radio. Furthermore, we believe the range of the data radio does not necessarily need to exactly match that of the wake-up radio. Ensuring the wake-up radio has at least the same or longer transmission range than that of the data radio will guarantee a node in data communications range can be interrupted and awoken. It is not vital to prevent all false positives. This is a view shared by other works in the literature, for example, IEEE 802.11 operates in a similar way. Although only a single radio is used, control traffic (e.g. RTS/CTS frames) is sent at a lower transmission rate, to improve signal recovery at the receiver, effectively increasing the transmission range of 802.11 control frames, over that of data frames.

4. RELATED WORK

To our knowledge, the work presented in this paper is the first implementation of a hybrid wireless communication system, utilizing ultra-low power radios in order to achieve the requirements of low power and low latency communication. However, the concept of utilizing ultra-low power radios as a wakeup control channel for higher power communications channels has attracted significant interest from the research community in recent years, particularly in the field of wireless sensor networks.

The concept was first discussed at the pioneering work at UC Berkeley as part of the now completed Picoradio project [5][6][7][8]. Results from simulation based analyses highlighted that systems based on this concept could experience significant benefits over schemes based on synchronized schedules alone. This work also alludes to the possibility of wakeup radios operating at a power consumption of 1uW or less.

More recently, Stankovic et al have published results of real world analyses documenting the feasibility of wake-up channels that are totally unpowered at the receiver [9]. Based upon similar concepts to passive RFID technology, this approach has been shown to provide up to 70% gains in energy efficiency over synchronized schedules in common WSN scenarios. The RF energy scavenging approach taken by the receiver was demonstrated to operate at up

to 10ft, and simulations of more advanced design which accumulate received RF energy over time, predict ranges of up to 100ft with latencies of 55ms. However, this cumulative nature may lead to a greater number of false positives.

Nogueira et al [10] published a system based on modulated backscatter technology. This allowed a device with a passive transmitter to alter the signal amplitude of a carrier wave between two other devices to convey data. They also described a promising RF wake-up circuitry with a running cost of only 12uA. However, their system only had a maximum range of five meters.

The STEM project also advocates the use of multiple radios where possible to facilitate wake-up channels in sensor networks, and its analysis and simulations indicate potential gains of up to two orders of magnitude over existing WSN approaches [11].

Other similar work has been carried out by the mobile computing community. Shih et al as part of the Wake On Wireless project demonstrated through extensive real-world studies the benefits of augmenting 802.11b wireless LAN networks with a relatively low power wake-up radio channel [12]. In these experiments, an 8mW radio transceiver was used in a 10:1 duty cycle allowing the device to operate for 115% longer than that of an unmodified device in typical office environments, with an additional latency of 200ms over 802.11b's existing power saving mechanisms.

Similarly, research at Intel and UCLA suggests that mobile devices equipped multiple radios such as 802.11b, Bluetooth and a CC1000 can be orchestrated to form a hierarchy of control channels for use in synchronization and discovery [13]. Experimental results indicate energy savings of up to 40x can be achieved in certain scenarios, but at the cost of significant additional latency.

Finally, there has also been much research effort placed into attempting to minimize energy consumption of sensor networks through optimized routing protocols and algorithms for synchronized sleeping schedules, though a detailed treatment of this research is beyond the scope of this paper [14][15].

5. CONCLUSIONS

Recently a new breed of interactive delay sensitive applications have emerged for energy constrained devices. This paper has argued why the traditional approach of sleep scheduling with long sleep cycles is not a suitable communication mechanism for such classes of applications, for reasons of high latency, high jitter communication, problems of resource discovery and synchronization problems caused clock skew.

We have also presented and evaluated a hybrid radio solution to support such applications that utilizes a low power radio to transmit addressable interrupts to other devices when communications with a higher power radio is desired. Our results have shown that although there is an additional cost of a secondary radio, the system was shown to be comparable with an 8 second wake up cycle followed by a 20ms active period, equivalent to a 0.25% duty cycle. These results have also shown that this scheme is not highly susceptible to interference of a congested channel, and that low power radios have comparable transmission ranges to modern higher power radios, despite being three orders of magnitude more efficient on the radio receiver. Furthermore, it has also been shown through practical implementation that this scheme is highly practical, and can be

implemented on existing ultra-low power radios and an 8-bit microcontroller with 32Khz clock, consuming only 100uA.

Finally, we believe having the ability to remotely wake nodes in a power efficient and effective manner goes some way toward addressing a number of issues in addition to latency, such as the overheads associated with node discovery and the effects of clock skew on sleep scheduling efficiency.

6. FUTURE WORK

We believe our findings so far have been very compelling, and so plan to continue to evaluate the potential benefits of the network interrupt system, predominantly in multi-hop scenarios. Firstly, we plan to investigate the use of cluster based addressing schemes towards enhancing the targeted wake up system. Secondly, we plan to evaluate the effects of network interrupts on current routing protocols and dynamic scheduling schemes. From the results of these investigations we expect to develop highly efficient devices and protocols to improve the overall support for real world low latency applications in challenged environments, and evaluate them via experimental analysis and field trials.

7. REFERENCES

- [1] C. Efstratiou, N. Daves, G. Kortuem, J. Finney, R. Hooper, M. Lowton, "Experiences of Designing and Deploying Intelligent Sensor Nodes to Monitor Hand-Arm Vibrations in the Field," Proceedings of Mobisys 2007, June 2007.
- [2] Chipcon CC2420 Radio, Texas Instruments, http://www.chipcon.com/files/CC2420_Data_Sheet_1_3.pdf
- [3] RF Solutions AM-HRR18 datasheet, http://www.rfsolutions.co.uk/acatalog/DS016-13_AM-HRR.pdf
- [4] NXP LCP213x, http://www.nxp.com/acrobat_download/datasheets/LPC2131_32_34_36_38_3.pdf
- [5] C. Guo, L. C. Zhong, and J. M. Rabaey, "Low Power Distributed MAC for Ad Hoc Sensor Radio Networks," Proceedings of GlobeCom 2001, November 2001.
- [6] J.M. Rabaey, M.J. Ammer, J.L. jr, da Silva, D. Patel and S. Roundy, "PicoRadio supports ad hoc ultra-low power wireless networking", IEEE Computer Magazine, July 2000, Vol. 33, Issue 7, pp. 42-48.
- [7] A. Willig, R. Shah, J. Rabaey and A. Wolisz, "Altruists in the PicoRadio sensor network", Proceedings of 4th IEEE International Workshop on Factory Communication Systems, 2002. pp. 175- 184
- [8] L. C. Zhong, R. Shah, C. Guo and J. Rabaey, "An Ultra-Low Power and Distributed Access Protocol for Broadband Wireless Sensor Networks", Proceedings of IEEE Broadband Wireless Summit, Las Vegas, May 2001.
- [9] L. Gu and J. Stankovic, "Radio-Triggered Wake-Up for Sensor Networks", Real-Time Systems Journal, special issue RTAS best papers, Vol. 29, No. 2-3, Mar. 2005.
- [10] J. Nogueira, R. Bohnke, O. Kramer, V. Wullich and J. Rebmann. "Topology aware - beaconless reactive wireless sensor network" Proc of Consumer Communications and Networking Conference, 2005.
- [11] C. Schurgers, V. Tsiatsis, S. Ganeriwal and M. Srivastava, "Optimizing Sensor Networks in the Energy-Latency-Density Design Space", IEEE Transactions on Mobile Computing, vol 1., no.1 January-March 2002.
- [12] E. Shih, P. Bahl and M.J. Sinclair, "Wake on wireless: an event driven energy saving strategy for battery operated devices". In Proceedings of the 8th Annual international Conference on Mobile Computing and Networking (MOBICOM 2002). September 23 – 28, 2002.
- [13] T. Pering, V. Raghunathan and R. Want, "Exploiting Radio Hierarchies for Power-Efficient Wireless Device Discovery and Connection Setup", In Proceedings of 18th International Conference on VLSI Design, pp. 774–779, 2005.
- [14] K. Akkaya and M. Younis, "A survey on routing protocols for wireless sensor networks", Elsevier Ad Hoc Network Journal, vol. 3, no. 3, pp. 325-349, 2005.
- [15] I. Demirkol, C. Ersoy and F. Alagoz, "MAC protocols for wireless sensor networks: a survey", IEEE Communications Magazine, Volume 44, Issue 4, April 2006 pp. 115 – 121.