Polymorphic radios: A new design paradigm for ultra-low power communication

Mohammad Rostami, Jeremy Gummeson, Ali Kiaghadi, Deepak Ganesan
College of Information and Computer Sciences
University of Massachusetts, Amherst, MA 01003
{mrostami,gummeson,dganesan}@cs.umass.edu,akiaghadi@umass.edu

ABSTRACT

Duty-cycling has emerged as the predominant method for optimizing power consumption of low-power radios, particularly for sensors that transmit sporadically in small bursts. But duty-cycling is a poor fit for applications involving high-rate sensor data from wearable sensors such as IMUs, microphones, and imagers that need to stream data to the cloud to execute sophisticated machine learning models.

We argue that there is significant room to optimize low-power radios if we can take advantage of channel dynamics in short-range settings. However, we face challenges in designing radios that are both efficient at power levels between μWs and mWs to take advantage of periods of good signal strength and nimble to deal with highly dynamic channels resulting from body movements. To achieve this, we propose radio polymorphism, a radio architecture with tightly integrated passive and active components that allows us to turn high channel dynamics to our advantage. We leverage passive modes in myriad ways within the network stack, from minimizing data transfer and control overheads to improving rate selection and enabling channel-aware opportunistic transmission. We instantiate our design in a full hardware-software prototype, Morpho, and demonstrate up to an order of improvement in efficiency across diverse scenarios and applications.

1 INTRODUCTION

Duty-cycling has emerged as the predominant method for optimizing power consumption of low-power radios. For example, Bluetooth LE is an optimization of the Bluetooth standard that enables rapid connection establishment, transmission of a short burst of information, and rapid disconnection. This rapid transition time makes it possible to mask the power consumed during active operation of the radio, which is milliwatts compared to microwatts in sleep mode. As a result, duty-cycled radios like Bluetooth LE and Zigbee have become the preferred choice for sensors that transmit sporadically in small bursts, for example, home temperature monitoring, location beacons, security alarms, humidity sensors, and other similar IoT devices.

But duty-cycling is insufficient for applications involving high-rate sensor data from IMUs, ECG, microphones, and imagers, that are used in wearable and tactile computing applications. The signals from these devices are noisy and complex which makes data interpretation a significant challenge [42]. To address this problem, we often need sophisticated machine learning techniques that are more complex than what we can execute locally and require computational resources in the cloud. The end result is a growing need for low-power radios that can support continuous streaming rather than transfer in short, intermittent bursts.

This trend has significant consequences from a power perspective. Normally, we would expect high-rate sensors to be the bottleneck in terms of power consumption but this has changed in recent years. For example, state-of-art low-power microphones, cameras, IMUs and ECG chips in the market consume between tens of microwatts and a few milliwatts for continuous sampling [9, 16, 27, 50]. But streaming communication has not kept pace with sensor developments — active mode power consumption of low-power radios is around ten milliwatts, which is an order of magnitude higher than the sensors. Continuous streaming of sensor data for real-time applications means that the radio needs to wake up frequently and cannot batch data before transmission. In these regimes, the prevailing wisdom of using duty-cycling to judiciously use the radio is ineffective.

Radio polymorphism: In this paper, we argue that there is significant room for optimization in the form of large gaps between received signal strength and receiver sensitivity. But these gaps often occur at extremely low power levels, and existing radio-level methods like transmit power control

SIGCOMM '18, August 20–25, 2018, Budapest, Hungary

ACM ISBN 978-1-4503-5567-4/18/08…$15.00
https://doi.org/10.1145/3230543.3230571

are too inefficient in these regimes. The problem is exacerbated by the highly dynamic nature of wireless channels from wearable peripherals due to body blockage and mobility. Thus, an ideal radio needs to be opportunistic and take advantage of gaps between signal strength and receive sensitivity while also being nimble and reacting quickly to channel degradation.

We argue for a new architectural paradigm for low-power radios, radio polymorphism, that tightly integrates active radio components like oscillators and active amplifiers with passive radio components like backscatter reflectors and envelope detectors. The key advantage of passive components is that they use extremely simple circuit components thereby allowing them to scale down power consumption to the microwatt regime. But compared to their active counterparts, passive radios suffer from lower signal strength, higher signal dynamics, and lower receive sensitivity. In other words, active radios are robust but inefficient whereas passive radios are efficient but fickle. The central challenge that we face is integrating active and passive radio components to accentuate their positives and mask their negatives.

To address this challenge, we take a step back and look holistically at integrating active and passive components while balancing energy-efficiency and robustness. Surprisingly, we find that there are several ways to approach the problem — passive components can be used for minimizing data transfer and control overheads, improving active bitrate selection, and enabling channel-aware transmissions, each leading to different ways of optimizing the overall system.

We put these ideas together in our instantiation of a polymorphic radio, Morpho, and extensively evaluate the benefits of the radio using benchmarks and trace-driven simulations. We also demonstrate the benefits of Morpho in two compelling application case studies: a) Morpho-enabled wearable eye tracking where we combine the radio with the iShadow eye tracker [28] to optimize gaze tracking performance without increasing the overall power budget, and b) Morpho-enabled audio streaming to optimize energy-efficiency without sacrificing audio quality.

Our work is distinct from a long line of work in multi-radio wireless communication (see §10). Most existing work does not specifically target the ultra-low power radio regime and generally looks at integrating WiFi, Bluetooth, LTE and other commercial radios at the MAC and transport layers. A small body of work has explored the integration of passive components in active radios but for specialized purposes like wakeup radios and power offload [17, 44]. There have been only preliminary efforts to design radios that truly integrate active and passive components [19, 43], and none that attempt to design the entire stack from hardware to application layers. Our work is a deep dive into this topic and unifies active and passive components across all layers of a wireless network stack.

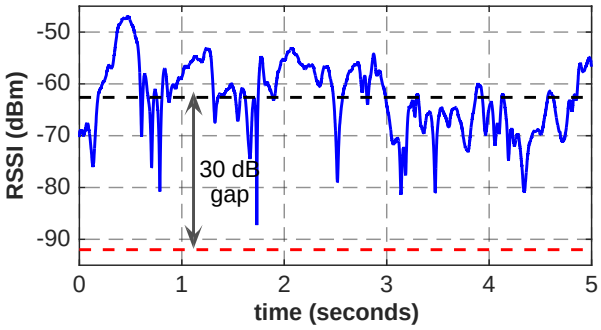

Figure 1: Gap between RSS and Rx sensitivity during short-range communication between a smartwatch and access point via Bluetooth @ 0dBm output power.

In summary, our work has several contributions:

- We present Morpho, a clean-slate re-design of ultra-low power radios that integrates active and passive components such that the modules operate in unison. Such a design combined with the ability to switch between modes in tens of microseconds allows us to optimize data transfer and control efficiency even under highly dynamic channel conditions.
- We show that Morpho provides 3.8× to 9× improvement in energy efficiency over active radios without compromising reliability under high channel dynamics that is typical in mobile and wearable scenarios.
- We show that Morpho can improve application-layer performance with two examples, a video-based wearable eye tracker where accuracy improves by 3× to 5× for a fixed energy budget, and an audio streaming application where energy efficiency improves by 5.8× - 10× while minimally sacrificing audio quality.

2 CASE FOR MORPHO

The primary opportunity to save power in low-power radios stems from the fact that communication is often short range i.e. within a few meters, whereas low-power radios are often provisioned to operate at ranges of a few tens of meters. This leads to a significant gap between received signal strength and receiver sensitivity that can be leveraged to save power.

Figure 1 illustrates this gap. The blue line corresponds to the RSS when a wrist-worn sensor is communicating with a proximate base-station via Bluetooth while a user is performing various gestures. The red dotted line shows the sensitivity of a typical Bluetooth receiver i.e. the lowest power level at which the receiver can detect an RF signal and demodulate data. In general, the receive sensitivity depends on thermal

noise given the channel bandwidth, the noise added by receiver electronics, and the required signal to noise ratio for the modulation scheme being used. In the case of Bluetooth, the thermal noise is -114 dBm for a 1 MHz channel [39] and the receive sensitivity is around -96 dBm at 1 Mbps [34]. As a result, in short range settings, the signal strength is often about 35 dB higher than receive sensitivity.

2.1 Leveraging the RSS-Sensitivity Gap

Thus, we often have a dramatic 30-40 dB gap between the received signal and the receiver sensitivity, but can we convert the opportunity into comparable power savings? There are two potential designs that can leverage this gap — transmit power adaptation and radio duty-cycling.

Transmit softly: Transmit power adaptation essentially involves reducing power consumption by dropping the output power of the transmitter so that the RSS becomes closer to the noise floor of the receiver. However, this does not lead to proportional power savings since the baseline operation of a low-power radio is already at a very low power level. For example, when a typical low-power radio transmits at 0dBm (i.e. output power of 1mW), the RF analog circuit consumes only around 5-10mW. If we wanted to reduce the output power from 0dBm to -30dBm (i.e. 1mW to 1μW) to take advantage of the gap, then we would need the RF analog circuit to operate at 10μW to achieve proportional power savings. But this is not possible due to the constant overheads of the active elements in a radio. In fact, the oscillator alone in a low-power radio consumes a few hundred microwatts, so power efficiency would be less than 1% when the output power is 1 μW [29]. Any other active elements like active mixers would only add to this overhead. Some of this inefficiency is apparent when we measure commercial low-power radios. For example, the Nordic nRF52840 BLE chip [34] draws 4.8mA when the transmit power level is 0 dBm and 2.3 mA at -40 dBm i.e. a 50% reduction in current draw for a four orders of magnitude reduction in transmit power. Thus, the fixed costs of a low-power radio swamp any gains that can be achieved by reducing transmit power.

Transmit rapidly: Radio duty-cycling involves transmitting at as high a bitrate as possible and saving energy by sleeping for longer. A higher speed PHY achieves lower power consumption (given that the same amount of data is transferred) since the radio-on time is reduced without changing transmit power. This is the approach used by virtually all low-power IoT radios. For example, BLE is typically configured to operate at either 1 Mbps or 2 Mbps to reduce power consumption.

But duty-cycling has two side-effects. The first is that the radio has no visibility into channel variations during radio-off periods. This means that mechanisms like rate adaptation are less effective in a duty-cycled radio since the channel may have changed since the last radio-on period. As a result, bitrates are often set to a fixed value in duty-cycled radios. The second is that constant overheads are significant for each wakeup. For example, a typical BLE radio goes through several stages during each wakeup cycle including MCU wakeup and shutdown, BLE protocol stack preparations and processing, and the radio on-off transitions [22]. The actual data transmission consumes only a fraction of the overall energy during each wakeup. These constant overheads can be masked if the messages are infrequent as is the case with BLE or when we can batch data to amortize the overheads. But they cannot be masked when streaming sensor data to the edge cloud. For example, real-time streaming of data from a microphone to an edge cloud (8 kHz sampling rate @ 16 bits/sample) via a 2 MHz Bluetooth radio would involve thousands of wakeups per second.

2.2 The Morpho Approach

We propose a new design paradigm that combines active radio architectures (i.e. RF oscillators, I/Q receivers, active mixers, power amplifiers, and low-noise amplifiers) with passive radio architectures (i.e. backscatter transmitters and envelope detectors). Such a design allows us to tackle the above issues in two ways. First, the constant overheads are a non-issue for passive radio architectures which do not have active components. Second, passive transmitters and receivers can operate in always-on mode and do not have to be duty-cycled since there is virtually no energy cost to using them. These advantages open up new possibilities in terms of how we can design low-power streaming radios.

But passive radios present a number of challenges that make it non-trivial to design an integrated active-passive architecture. In the case of a passive transmitter (i.e. backscatter), the main issue is substantially higher path-loss. Since the backscatter signal has to traverse the forward path and the reverse path, the attenuation is exponentially greater than an active radio where the signal only needs to traverse the forward path. In effect, this is a double-whammy since the signal average is considerably lower than an active radio, and the signal dynamics is a considerably exaggerated version as that for an active radio. The challenge is not limited to the transmitter — a passive receiver (envelope detector) also presents problems since its sensitivity is often considerably lower than an active receiver. Thus, when we integrate these two vastly different radio architectures, we need to carefully consider how we accentuate their positives and mask their idiosyncracies in-order to improve performance.

Morpho presents a unification of active and passive modules into a single radio that transparently switches across these modules to optimize energy-efficiency without losing

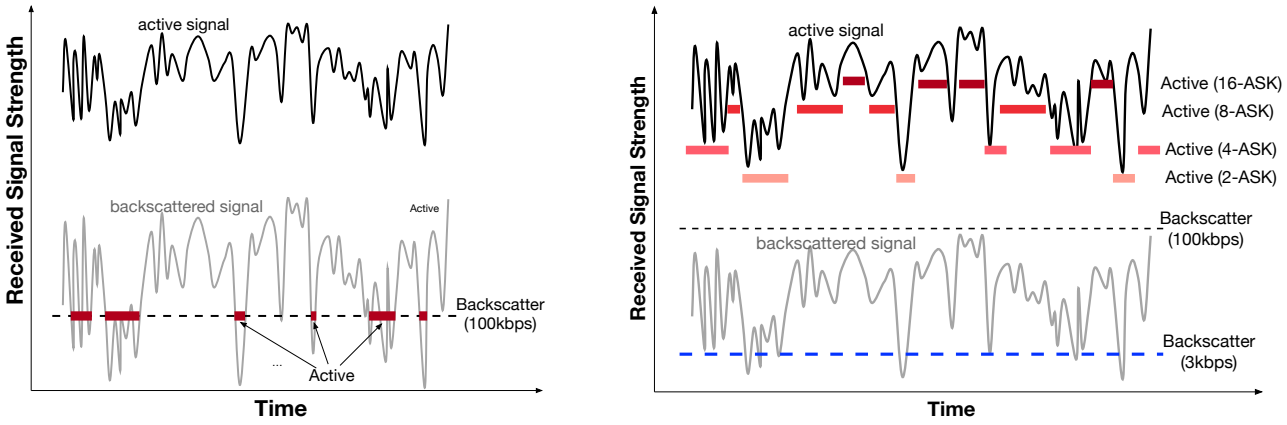

(a) Active-assisted Backscatter

(b) Backscatter-assisted Active

Figure 2: We illustrate the main ideas in Morpho for uplink transmission. Active-assisted Backscatter shows the case where backscatter works at sufficiently high bitrate most of the time, so the active radio is only needed as backup when backscatter fails. Backscatter-assisted Active shows the case where backscatter cannot support a high bitrate for data, but it can be used at very low bitrate for zero-power channel measurements. This enhanced visibility allows us to choose the best bitrate and time to use the active radio, thereby improving efficiency.

robustness. The application is agnostic to the manner in which Morpho switches between modules even when it includes rapid transitions needed to handle highly dynamic channels that are typical in mobile situations. Our vision is to enable a deep stack integration, where the physical layer, protocol layer, and application layer are all re-architected to squeeze the most out of opportunities to use passive radio modes without sacrificing the intrinsic robustness of low-power active radios.

3 DESIGN RATIONALE & KEY INSIGHTS

How should we combine active and passive modes to design a unified radio? When considering the answer to this question, we found that there are two distinct approaches to unify the two modes. We describe these approaches in this section.

3.1 Active-Assisted Passive

The first method for combining active and passive modes is somewhat evident — use the passive mode whenever available for data transfer since it has better energy-efficiency and use the active mode to smooth out periods when the passive mode is flaky. We refer to this mode as *active-assisted passive* since passive is the preferred data transfer mode and active is backing it up by filling connectivity gaps.

Let us first consider the case of a peripheral (e.g. IoT sensor) transmitting to a central station (e.g. access point) via a polymorphic radio. Here, the choice is between transmitting via a (passive) backscatter transmitter or via an (active) I/Q transmitter. This scenario is shown in Figure 2(a) — the bold line is the active signal strength and the grey line is the backscatter version of the same signal. The dotted horizontal line represents the receive sensitivity at a desired data rate, 100kbps in this example. The backscatter signal has lower RSS and higher signal variation compared to the active signal, but despite this the signal is mostly above the receiver sensitivity and can be decoded. In this case, backscatter can be used for data transmission most of the time with active being used as backup whenever the backscatter signal strength goes below the receive sensitivity.

We note that there are several parameters that can be tuned to change the operating region shown in the figure. The first is the carrier signal power from the central station. The second is the acceptable bitrate — the receive sensitivity line in Figure 2(a) can be lowered if the acceptable bitrate for the passive mode is less than the stated 100 kbps figure.

This approach is equally applicable to the case of a peripheral receiving data from the central station. Here, the peripheral has to decide whether to use the (passive) envelope detector or to use an (active) I/Q detector. In an active-assisted passive approach, the peripheral uses the passive mode whenever the received signal is strong enough to use the passive envelope detector as the primary receiver, and the active mode kicks in when the signal falls below the sensitivity of the passive receiver (at the desired bitrate).

3.2 Passive-assisted Active

While active-assisted passive is ideal under conditions where the passive modes offer sufficient throughput, there are often

conditions where bitrate offered by the passive mode is too low. For example, when active RSS is -65 dBm, the corresponding active bitrate for 16-QAM modulation is 4 Mbps whereas the corresponding backscatter bitrate would be ~5 kbps. Thus, there are often scenarios where the difference in throughput between active and passive mode is too high for active-assisted passive to be practical. The question we ask is whether we can still leverage the passive modes to improve performance in these situations.

Our main insight is that even though backscatter may be impractical for supporting data transfer, it can still be useful for *channel measurement* at extremely low bitrates.

Let us again consider the case of a peripheral transmitting to a central station via a polymorphic radio. In order to measure the channel, a short training sequence of a few bits can be transmitted, and the RSS estimated by obtaining the correlation between the received signal and the training sequence. Since the training sequence can be as short as a few bits, the backscatter bitrate can be as low as a few kilobits/second which can allow it to operate at longer distances (e.g. 75 m @ 2.9 kbps [49]). Figure 2(b) illustrates this idea — the grey line corresponding to backscatter RSS is below the receive sensitivity when operating at 100 kbps but mostly higher than the sensitivity when transmitting at 3 kbps.

The central benefit of being able to use backscatter for channel measurement is enhanced visibility into the channel at near-zero power consumption. The additional visibility allows us to be more judicious about use of the active radio in two ways: a) we can select the best bitrate for the active radio even after a long sleep gap, and b) we can select the best times to wakeup the active radio by choosing times when the RSS is strongest. The figure shows these advantages — by leveraging backscatter for channel visibility, we can choose the best active bitrates (the line-segments) and the best times for active transmission (the peaks).

We note that this approach does not directly translate to the scenario where the peripheral is receiving data from the central station. This is because an envelope detector does not provide signal strength information, so cannot be used for channel measurement. Therefore, we use passive-assisted active solely for uplink transmission from the peripheral to the central station.

4 MORPHO PHY LAYER

The main challenge at the PHY layer is how to seamlessly integrate passive and active radio components so that they can transparently switch between various modes without the application perceiving the switching behavior. To accomplish this, we design the Morpho hardware to enable rapid and seamless switching.

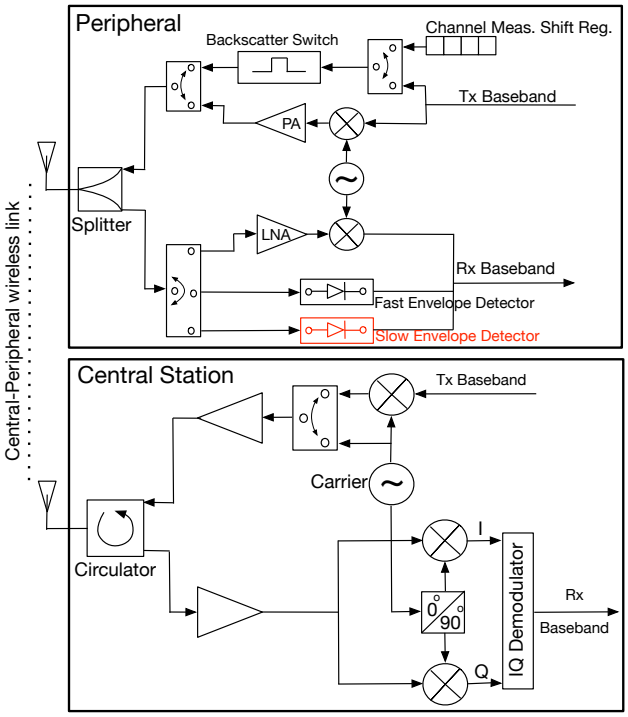

Figure 3: The building blocks of Morpho hardware.

To illustrate the need for a new design, we start with a strawman solution for an active-passive radio that simply connects an active radio like BLE together with a passive radio (e.g. WISP), and switches between these two as needed (similar to BLISP [19]).

Such a design is inefficient due to the lack of configurability. Virtually all radio ICs have the TX and RX components tied together, and do not provide us the freedom to mix-and-match different possibilities. For example, low-power active radios like BLE turn on both the TX and RX components when they switch on from sleep mode since they assume that active-mode ACKs will follow the data packet. In addition, most commercial radios incur setup delays upon receiving a command, which makes it difficult to rapidly switch between the different modes. In contrast, our goal is to have the freedom to rapidly switch between all four combinations of passive rx/tx and active rx/tx.

4.1 Morpho Sensor Architecture

Morpho is designed for one-hop asymmetric settings where the peripheral (e.g. IoT sensor or mobile accessory) is resource-constrained whereas the central station (e.g. access point or edge cloud) is resource-rich. We first describe the radio architecture on the peripheral shown in Figure 3 (upper block).

Morpho TX: The Morpho transmitter is equipped with an ultra-low power RF oscillator that shares the antenna path

with a backscatter transmitter allowing us to turn on the oscillator and send data in an on-demand manner without incurring additional overheads.

The figure also shows the channel measurement circuit for the backscatter-assisted active mode. Channel measurement using low-rate backscatter can be implemented as a simple shift register in hardware since a fixed set of bits are backscattered each time to measure the channel. This makes it extremely low power (similar to an RFID tag) since it avoids the overhead of waking up the MCU.

Morpho RX: On the receiver end, the radio has a switch between an envelope detector receiver and an active ASK receiver (which shares the oscillator used by active transmitter). One notable design issue that we encountered was that the envelope detector receiver needs to be tuned to specific bitrates (and consequently operating ranges) — high bitrates need low RC constants and long ranges need high RC constants. This is different from a backscatter transmitter which can transmit at different rates simply by toggling the RF transistor at different speeds. Thus, we were presented with a tradeoff between bitrate and range.

We therefore used two envelope detectors — one specifically tuned for data transfer and the second tuned for longer range and lower rate operation. This is shown in Figure 3, where a second envelope detector (in red) is tailored for low rate control messages where range is more important than rate. A significant sensitivity gap can be expected between these two detectors — for example, a state-of-art detector for rates of 100 kbps has a receive sensitivity of -50 dBm whereas a detector for low rates of a few kbps has a receive sensitivity of -68 dBm [52].

4.2 Morpho Central Station Architecture

The Morpho central station (or base station) is a more power-hungry system since it needs to generate the carrier whenever the peripheral is operating in backscatter mode. The architecture of the central station resembles that of a typical backscatter reader but with the difference that it can switch between being generating a carrier when needed to support backscatter at the peripheral and operating as a standard active receiver when the peripheral is transmitting in active mode. The central station needs methods to deal with self-interference when generating the carrier for backscattering from the peripheral since the carrier can overwhelm the weak backscattered response from the peripheral. There are many approaches to perform carrier cancellation [17]; Figure 3 (lower block) shows an approach that relies on a circulator [39].

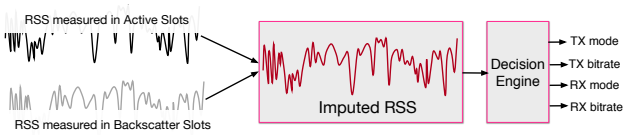

Figure 4: Prediction model for deciding whether to use active or passive mode. The figure shows input coming from the backscatter and active RSSI. Since only one of them is used in a given slot, we have to impute readings before feeding the values to the predictor. Since backscatter and active RSSI are roughly linearly related, we can impute using a simple function.

5 MORPHO MAC LAYER

We now have a radio that can rapidly switch between active and passive modules but we now need to orchestrate these components to optimize robustness and energy-efficiency. Morpho is a master – slave system where the central station controls the operation of the peripheral and makes decisions regarding the TX and RX modes of the peripheral. Thus, the MAC layer is based on a simple TDMA protocol that is driven by the central station similar to other backscatter-based protocols.

We first describe the decision engine is responsible for tracking channel dynamics and deciding between the various passive and active modes and then the MAC layer protocol that provides the rubric for switching between the modes.

5.1 Decision Engine

The decision engine is responsible for tracking channel dynamics and deciding between the various passive and active modes. The decision engine has two key components as shown in Figure 4.

Imputation of active and backscatter RSS: A unique feature of the Morpho decision engine is that it has visibility into the channel even when the active radio is not being used. However, since the active and backscatter radios are interleaved, we have RSS for only one of these radio modes for each slot and have to impute the missing data. The imputation function leverages the fact that Active Tx incurs only one-way pathloss whereas Backscatter Tx incurs two-way pathloss. There are also several constant offsets due to carrier self-interference and transmit power level but these are known a priori and can be accounted for.

To deal with noise, we impute not only using the RSS in a particular slot but also the RSS of previous N-1 slots. In slots for which we have no information i.e. when backscatter or active fails entirely, we use a pre-defined RSS that is below the detection threshold of the receiver. The output of the

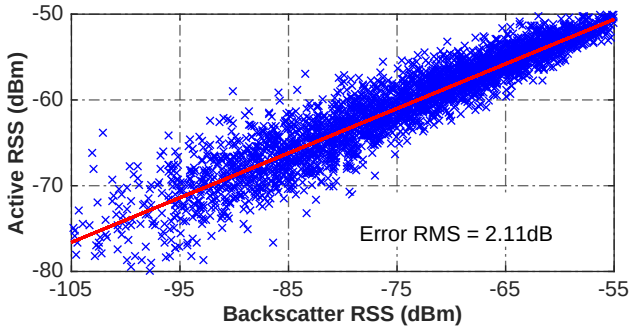

Figure 5: A linear relationship is observed between the backscatter RSS and active RSS.

imputation is the smoothed active and backscatter RSS for the past N slots.

Figure 5 shows the relationship between Backscatter and Active RSS for communication between a smartwatch to base-station during normal movements. We can see that the relationship is generally quite linear (in dB) since links are symmetric at short ranges and the forward and reverse path are typically the same. However, there is more measurement error when we are near the receive sensitivity of our measurement infrastructure since the noise levels are higher.

Prediction engine: Given the imputed signal, the prediction engine makes the following decisions: a) for data messages, it needs to decide between backscatter data, active data, and no data (i.e. just measure channel via backscatter), and b) for control messages, it needs to decide between passive high-rate receiver, passive long-range receiver, and active receiver.

The prediction engine first looks at the most recent backscatter RSS measurement and if it is above the threshold for transmitting data via backscatter at the desired bitrate, then it decides to send immediately via backscatter. Since backscatter has very low cost, there is nothing to be gained by waiting for a better RSS. If backscatter for data is not a viable option, it needs to decide whether to send immediately with one of the active modes or wait for a better RSS (within an application-defined latency window).

In order to do so, the prediction engine needs to look at the trends in RSS variations by using all measured RSS samples within the several past time windows. Let us define T to be the time (in slots) from the current slot until a better RSS will appear for the first time. We compute the probability distribution of T, given two parameters: the current RSS, and the current RSS slope. We define segments in the 2-D space of RSS values and slopes and obtain the distribution of T for every cell. Given the distribution, our goal is basically to determine whether at the current slot there is a high probability of having a better RSS before the window ends. We define t to be the number of slots until the window ends. Therefore, we must look at:

$$P = \text{Prob.}\{T \leq t \mid \text{current RSS}, \text{current slope}\},$$

and if it is below a threshold (80% in our implementation), the decision is to send via active at the bitrate determined by the imputation process, and if not, the decision is to wait. Note that waiting is the same as measuring the channel via low-rate backscatter, so we continue to have visibility into the channel.

5.2 MAC Layer Protocol

At the protocol level, we design an integrated MAC layer that is able to switch between four modes — Active TX, Active RX, Backscatter TX and Passive RX — as and when needed based on the results of the decision engine.

Since Morpho is a master – slave system, the central node needs to inform the peripheral regarding which mode to use. To enable this, every slot is partitioned to a small control sub-slot during which the central node sends the control commands to the peripheral, and a bigger data sub-slot for sending or receiving data bits.

Figure 6 shows a sequence of slots in the case of uplink data transfer from the peripheral to the central station. During the control sub-slot, the central station sends a command to the peripheral which provides information about: a) the uplink mode and bitrate, and b) the downlink mode and bitrate, and c) an ACK for the data transfer that occurred in the previous slot. Each of these is only a few bits, so the overhead is small. The figure shows cases where the peripheral transmits uplink data using backscatter (top), performs a backscatter channel measurement using shift register (middle), and transmits uplink data using active mode (bottom).

To ensure robustness, we fallback to the active modes when failures occur in the passive modes. In particular, the control channel needs to be as reliable as possible since it drives the behavior of the system. Therefore, whenever the passive receiver fails in the control sub-slot, the peripheral switches to the active receiver in the subsequent control sub-slot. Since the control sub-slot is generally much smaller than the data sub-slot, overall power efficiency is not significantly compromised by using an active receiver for control. If backscatter transmission in the data sub-slot fails, then the decision of whether to retransmit or switch to active mode is provided by the central station in the next control sub-slot.

6 RE-THINKING APPLICATION DESIGN

Unlike typical low-power radios which offer some limited capability to tune power consumption, Morpho is unique in that its power consumption can vary by three orders of magnitude depending on whether active or passive modes are

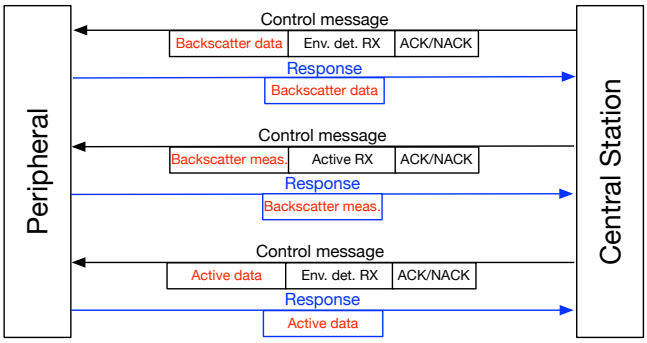

Figure 6: The Morpho MAC layer. Three exchanges are shown. (top) central station commands peripheral to use backscatter uplink in next data sub-slot and envelope detector downlink in next control sub-slot. (middle) central station asks peripheral to use active uplink, (bottom) peripheral asked to use low-rate backscatter to measure channel.

being used. We now describe two ways in which applications can leverage this power gap to improve performance.

Quality-Power tradeoffs in audio streaming: The most straightforward way in which an application can use Morpho is to use the passive radio whenever it provides sufficient bandwidth. Morpho then uses the passive mode whenever RSS is high enough to support the required throughput, and if that is not possible, it tries to use the passive mode for channel measurement. If the passive mode does not work at all, it exclusively uses the active radio as a traditional low-power radio. To illustrate this approach, we consider audio streaming using an application like Skype or Pandora which can leverage Morpho to tradeoff application performance for significant gains in power consumption. Audio streaming typically operates at low rates of 32–64kbps for speech and 128kbps for audio, and such bandwidth is frequently achievable using passive communication at short range. This gives Morpho the opportunity to leverage passive communication aggressively and tradeoff a small reduction in audio perception quality for substantial power gains.

Eye tracking with adaptive sampling: Morpho can also be used in concert with the application — as the radio adapts to dynamics and adjusts its operating point along the active–passive spectrum, the application layer can also adjust its computation and sensing decisions to leverage the ultra-low power operation in passive modes. This adds a new dimension to how we opportunistically use cloud and local resources to improve application performance.

We illustrate these advantages with a case study involving a wearable eye tracker [28] that uses sparse sampling to sample pixels from an imager, and uses a neural network to

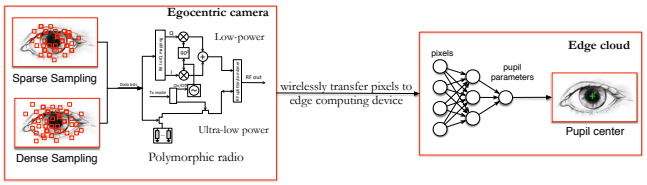

Figure 7: Eye tracking with cloud offload. Depending on which radio is available, the eye tracker either densely samples the pixels or sparsely samples them to transmit to the edge cloud for processing. Thus, it adapts accuracy based on channel dynamics.

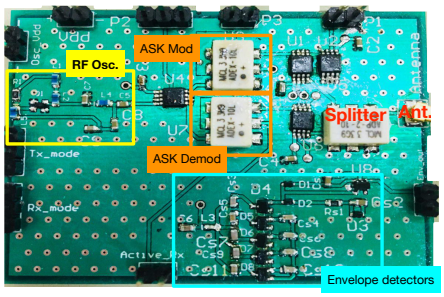

Figure 8: Morpho prototype

compute gaze parameters as shown in Figure 7. In this example, a polymorphic radio can be used to adapt the amount of energy used for sensing vs communication. If more energy is expended for communication via an active radio, then the eye tracker can reduce the pixels sampled and transmit fewer pixels to reduce the overhead of using the the active radio. If less energy is consumed for communication via backscatter, then the camera can be sampled more densely, and these pixels transmitted cheaply to an edge cloud for processing. Thus, by working in concert with the radio, the application layer can improve its accuracy.

7 IMPLEMENTATION

The main challenge that we faced in implementing Morpho is the complexity of building a radio and its protocol stack from the ground up. Our design has many non-traditional requirements including: a) power adaptation from microwatts to milliwatts, b) configurability to enable arbitrary combinations of active and passive modes with tiny switching overheads, and c) multiple backscatter transmission subsystems and multiple envelope detectors tuned for different purposes. The combination of these meant that the use of off-the-shelf components and transceivers were essentially off the table.

Trace	Description
$\mathbf{T}_1$: Wrist IMU	Streaming data from a Smartwatch Inertial Measurement Unit (IMU) to central station for gesture recognition. Streaming 100 samples/second from a nine-axis IMU. Medium average throughput of 10 kbps.
$\mathbf{T}_2$: Lapel Audio	Streaming audio from a lapel sensor for dialog-based applications. Streaming audio at 4 kHz. High average throughput of 32 kbps.
$\mathbf{T}_3$: Eyeglass camera	Streaming video from low power camera on an eyeglass for first-person vision applications. Streaming video at 30 (sub)frames per second - every (sub)frame consists of 800 pixels. Average throughput of 240 kbps.
$\mathbf{T}_4$: Audio download	Same scenario as $\mathbf{T}_2$ but audio is streamed from the central device to peripheral (e.g. music).

Table 1: Description of experimental traces. In all cases, we assume that the data is streamed roughly sample-by-sample with a low latency of 30ms. In each case, we collect simultaneous channel information in both passive and active modes, allowing us to compare strategies.

7.1 Morpho prototype

One issue we encountered was that ultra-low power RF oscillators with tight sleep-active transition times were not available as stand-alone components for PCB-level integration. Hence, we custom-designed a Collpitt LC oscillator and integrated it with the backscatter radio (RF oscillator block in Figure 3). The oscillator was first designed in ADS simulation environment in order to tune its LC components as well as the lengths of micro-strip tracks to the right frequency and output power, then implemented on a PCB with an NXP BFU690F NPN RF transistor [35]. The Collpitt oscillator that we designed has an output power of +1.1 dBm and wakeup time of 25-35 μs, and generates a 910MHz carrier.

For completeness, we also mention other components of the design. We use a ADG902 SPST RF switch [3] is used as backscatter switch, and HSMS-285C Schottky diodes [6] are used in the envelope detectors. An ADEX-10L+ passive mixer [31] is used for implementing higher order ASK modulations to change the active bitrate as needed. Also, ADG919 SPDT RF switches [4] are used in order to multiplex between the [Tx/Rx $\rightleftarrows$ antenna] paths and to switch between active and passive modes. Finally, an ADP-2-10+ RF power splitter [32] is used to split the Tx and Rx paths to the antenna, and we use a W1910 1dBi small whip antenna [40] as our antenna for Morpho prototype.

On the digital side, the packetizer, MAC layer controller, and the low bit rate sequence generator used for measurement is implemented externally on a AGLN250 low power FPGA development board [30], which connects to the prototype via the connectors shown in Figure 8.

7.2 Base station implementation

Our base-station is built based on a X300 USRP [13] operating at +30 dBm carrier. Since the base-station must be able to work in both backscatter and active modes, we use an ADG902 evaluation board [3] to turn on and off the carrier. The entire decision engine and data decoding stack is implemented inside a Mac mini computer that is connected to the USRP and to the ADG902 switch. The UBX-40 daughter board has -100 dBm noise level for 1 MHz bandwidth.

On the software side, we run our switching and control tasks in MATLAB, and transfer raw IQ samples from USRP to the MATLAB environment using a TCP socket. There is a 200 μs-300 μs latency in the connection, which has a small effect on overall performance. Within MATLAB, we implemented several software modules including backscatter RSSI measurement, active RSSI measurement, ASK demodulation, data imputation, and prediction engine.

8 EVALUATION

Our evaluation uses a combination of trace-driven evaluation and real experiments. Since we need to compare many different communication strategies on the same underlying channel dynamics, our benchmark and comparison results are evaluated on traces where we simultaneously collected data from the active and passive radios. The application studies are based on a live implementation.

8.1 Hardware micro-benchmarks

We start with hardware micro-benchmarks before showing overall performance results. Table 2 benchmarks the performance of Morpho at the hardware level. The most important optimization is the ability to switch within tens of microseconds between passive and active modes allowing us to respond swiftly to channel dynamics. The result also shows that we operate at extremely low power levels while measuring the channel via backscatter. We also see that our Colpitt oscillator performs efficiently at output power levels typically used by low-power radios.

The table also shows the benefits of using both a high-rate but short-range detector and a long-range but low-rate detector. The fast detector supports good bitrates of 32 kbps to 128 kbps but can only operate at high RSS levels of roughly

Component	Performance
Mode Switching - Latency	30 μs
Mode Switching - Power	5.2 mW
Active mode	5.2 mW @ 1.1 dBm
Backscatter TX (data)	50 μW
Backscatter TX (measurement)	10 μW
Passive RX (env. detector)	10 μW
— Env. Detector 1	< 32kbps, -28dBm sens.
— Env. Detector 2	32–128kbps, -20dBm sens.

Table 2: Morpho micro-benchmarks showing low-power operation and tight switching latency.

-20 dBm whereas the slow detector operates down to -28 dBm but only supports bitrates of up to 32 kbps. The combined detector covers the superset of the two receivers.

8.2 Morpho vs. active and passive radios

In this section, we validate our claim is that Morpho provides the robustness of active radios and the efficiency of passive radios. To do so, we compare Morpho against a duty-cycled fully active radio and against a fully passive radio.

Data traces: In-order to perform a fair comparison between the three schemes under the same channel conditions, we obtain four traces corresponding to exemplar applications that involve high rate communication from or to wearable device (as shown in Table 1). Of these, the first three are upload-intensive and increase in data rate from T_1 to T_3 and the last one (T_4) is download-intensive. For each of these traces, we collect simultaneous channel information in both passive and active modes, allowing us to compare strategies.

We use a scripted procedure to collect these traces. We first divided the whole experimental area which is a large 7m × 6m room to three sub-regions based on the distance to the reader: Short-distance (~ two meters), Medium-distance (~ four meters), and Long-distance (~ six to seven meters). Then, we designated 10 locations in each sub-region in order to cover the space of distances between the Morpho node and the base station. For each trace, we placed Morpho at the appropriate spot on the body and walked between the predefined locations while spending 30 seconds at each location. We also scripted a set of natural gestures to be performed at each location including natural movements of the hand while picking up an object, moving the head and hands normally while speaking, and turning the body.

In-order to evaluate the three methods over these traces, we implemented the complete Morpho MAC layer in MATLAB with parameters obtained from hardware micro-benchmarks. To avoid differences across hardware platforms, we assume

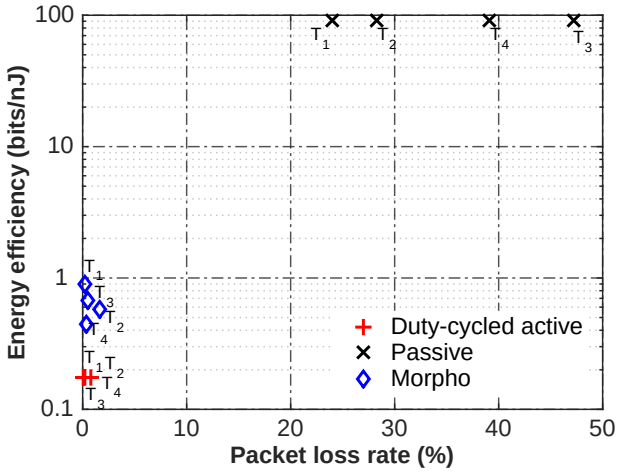

Figure 9: Energy Efficiency v.s. Packet loss rate of Passive, Active, and Morpho for all traces.

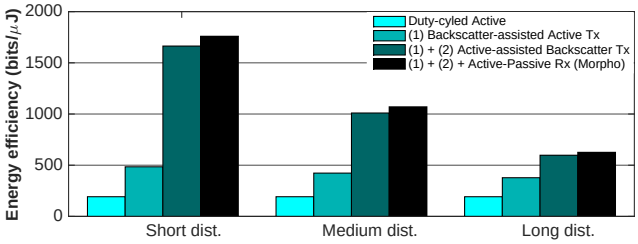

Figure 10: The effect of distance on the contribution of each building block towards overall energy efficiency (for Wrist IMU trace, T_1).

that active and passive are executing over the Morpho hardware prototype. We assume a 1 ms TDMA slot size (roughly the size of an EPC Gen 2 slot [53]).

Overall performance: Figure 9 shows the packet loss rate vs. power consumption of the three methods. Duty-cycled Active and Morpho methods have very low loss rates (0.1%-%1.5), whereas Fully Passive has 25% – 50% loss rate. Note that Morpho has marginally higher loss rate than active because of occasional decision engine errors such as choosing passive rather than active, or too high a bitrate for active communication. In terms of energy-efficiency, Morpho is between 2.5× — 5× more efficient than duty-cycled active radios depending on the specific trace. These results validate that Morpho provides a balance of robustness and efficiency by intelligently using the two radio modes.

We now breakdown the above results in several ways to better understand the contribution of various building blocks of Morpho to the overall performance.

Results by distance: Morpho works best at shorter distances where passive modes can be heavily relied upon for

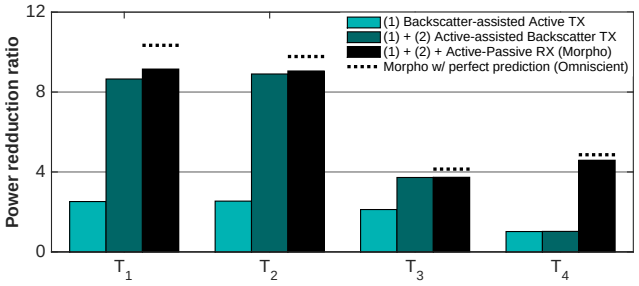

Figure 11: Gains due to each of the building blocks of Morpho for short-range communication. Gains are computed relative to Duty-cycled Active on that particular trace.

transmission and reception of data. Figure 10 illustrates this effect in the instance of the Wrist IMU scenario (T_2). We see that largest gains for Morpho come at short distances where backscatter can be heavily used for data transfer. At this distance, Morpho is about 9.1× more efficient than an active-only approach. At larger distances, the benefits are roughly equal for active-assisted backscatter and backscatter-assisted active and the improvement is about 3.3× — 5.6×. Since T_2 primarily involves data upload from the peripheral to the central station, the contribution of the active-passive receiver is low.

Breaking down the benefits: We now look at how the building blocks of Morpho contribute to the overall performance. Figure 11 compares the contribution of the different building blocks of Morpho normalized against the power consumption of duty-cycled active for that specific trace. We zoom into the short-range part of all traces since all the building blocks of Morpho work together in this regime. The figure shows the contribution of the three main innovations in Morpho — active-assisted backscatter (i.e. using backscatter for data), backscatter-assisted active (i.e. using backscatter for measurement), and use of active-passive receivers for control. For reference, we also show an omniscient version of our decision engine that can predict the perfect policy.

The contribution of different building blocks varies across the traces. Let us first look at the upload-intensive traces T_1 — T_3. Backscatter-assisted Active provides a steady benefit of roughly 2.5× across these traces. In contrast, the benefit from Active-assisted Backscatter is roughly 6× for T_1 and T_2 whereas it is only 1.6× for T_3. This is because the data rate is lower for T_1 and T_2, hence there is more time to wait until channel conditions improve such that the passive mode can be used. Let us now look at the download-intensive trace T_4. Here, almost all of the 4.5× improvement comes from active-passive receiver rather than from transmitter optimizations.

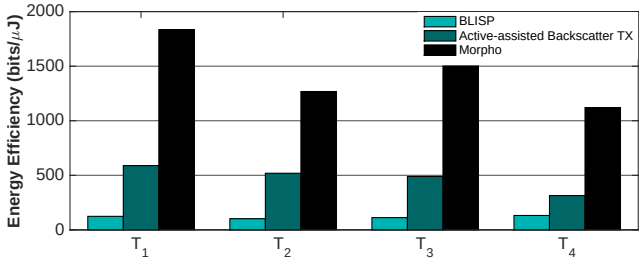

Figure 12: Comparison of Morpho against BLISP [19].

We also see that Morpho performs close to the omniscient scheme that has perfect knowledge of channel conditions.

Impact of rapid switching: One of the advantages that Morpho provides is the ability to switch rapidly in a manner transparent to upper layers. We now look at the benefits of rapid switching for the different traces.

Trace		T_1	T_2	T_3	T_4
Data Slot	Active	33%	44%	64%	53%
	Passive	67%	56%	36%	47%
Control Slot	Active	25%	36%	47%	34%
	Passive	75%	64%	53%	46%
Switch rate (per second)		11.1	15.4	19.3	21.4

Table 3: Percentage of time spent in active and passive modes in data/control slots, and aggregate switching rate in each trace)

Table 3 shows the rate and fraction of switching between different modes for the traces. We see that switches between the two modes are frequent and occur roughly 10–20 times per second. This validates the need for a radio that can rapidly transition between modes to adapt to a dynamic channel in-order to minimize energy overheads while providing a unified abstraction of a single radio to upper layers.

To further illustrate the benefits of fast switching, we contrast Morpho against BLISP [19], which combines a BLE active radio with a WISP passive radio [46]. BLISP relies on an algorithm similar to active-assisted backscatter i.e. it uses backscatter mode when available and active mode when backscatter fails. Since BLISP relies on commodity radios, it has high switching latency and incurs more overhead for transition between modes. We empirically measured the switching latency and power from deep sleep to active mode for BLE (560μs & 11.3mW respectively), and use these parameters for the BLISP comparison.

Figure 12 shows that Morpho is 8.5× to 14.8× more efficient than BLISP across the traces, and even the active-assisted backscatter component within Morpho is, by itself, between 2.4× – 5.1× more efficient. This is because Morpho

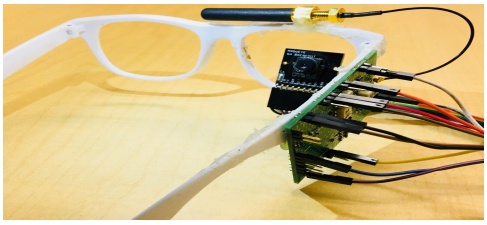

Figure 13: Prototype of eye tracker with Morpho

is far more nimble than BLISP with tight switching capability, but also because it has several additional design elements including backscatter-assisted active and active-passive control optimization.

Method	Energy Efficiency (bits/μJ)			
	T_1	T_2	T_3	T_4
Baseline (use prev. slot RSSI)	685	545	321	244
Morpho Prediction	1092	968	535	556

Table 4: Benefits of prediction.

Benefits of predicting active channel: We now look at the benefit of our prediction scheme against a baseline method that assumes that the RSS for the current slot is the same as the RSS for the previous slot. Table 4 shows that our prediction method improves energy efficiency of Morpho by roughly two times over the naïve prediction method which cannot deal with a dynamic channel.

8.3 Application-layer Performance

We now consider two applications that leverage Morpho and evaluate how the radio can improve their performance. The results in this section are based on a full hardware-software integration to enable live experimentation.

Eye tracking: Here, we illustrate the benefits of Morpho for an eye tracker whose sampling decisions are varied based on the current power consumed by the radio. We integrated Morpho with an eye tracker as shown in Figure 13.

While a complete description of the eye tracking mechanism and hardware can be found in [28], we describe salient details to understand the results. Briefly, the eye tracker is able to sample the imager at different resolutions and extract user gaze location by running neural network models trained for different sampling patterns. Clearly, accuracy depends on the sampling resolution and varies from a gaze error of 10—15 pixels at 77 pixels/frame to 0—3 pixels at 1984 pixels/frame, where error is the euclidean distance between actual and predicted gaze location.

For this evaluation, we assume that the eye tracker device has a fixed power budget of 100μW for every gaze location update (at 30Hz frame rate). The tracker adjusts the sampling rate depending on available budget — when communication costs more, it samples less and vice-versa. We follow the same procedure that we used to collect the traces in Table 1 i.e. we move around a room in a scripted manner with the main difference being that we were running a live version of the eye tracker.

Method	Gaze Err (all dist.)	Gaze Err (short dist.)
Active	9.1 ± 5.7	8.8 ± 6.0
Backscatter	17.8 ± 8.6	10.6 ± 7.4
Morpho	3.6 ± 3.4	2.7 ± 2.7

Table 5: Optimizing eye tracking with Morpho.

Table 5 shows that Morpho reduces gaze error by about 3× over an active-only approach and 4–5× over a backscatter-only approach (both at short distance and across all distances). The improvements occur because Morpho is able to use energy saved in communication on sensing, thereby transmitting more samples and improving accuracy.

Method	Short distance		All distance	
	MOS	bits/nJ	MOS	bits/nJ
Active	4.4	0.17	4.3	0.17
Backscatter	2.85	100	2.51	100
Morpho	4.1	1.74	4.0	0.98

Table 6: Audio voice quality over Morpho versus duty-cycled active and backscatter-only.

Voice audio streaming: We now look at streaming voice audio over Morpho (upload from peripheral to central station). For this experiment, we attached Morpho to a shirt (convenient location for microphone) and transmitted a stored audio stream via the different modes. We use typical audio streaming parameters (40kbps rate, 125 byte audio packets, and 30ms latency). We followed the same procedure as the eye tracking example in terms of moving across different locations in the room while the data was being transmitted. We then computed the Mean Opinion Score (MOS) of the audio stream [41].

Table 6 shows that MOS score is marginally lower than for the active radio (primarily because of bit rate changes when switching between active and passive modes) but energy efficiency is four times higher across all distances and an order of magnitude higher at short distance. Backscatter has considerably lower MOS score but it has very high energy-efficiency. Thus, Morpho is able to take advantage of the channel to improve energy efficiency without significant impact on application-level performance.

9 DISCUSSION AND LIMITATIONS

We briefly discuss some additional issues that we did not cover in the rest of this paper.

RF tuning and hardware optimization: We expect that the performance of Morpho can be increased substantially with better RF optimization. Other work has reported tens of meters backscatter range and higher sensitivity passive detectors [18, 33, 49, 52]. Such improvements can extend our techniques to larger sized areas (e.g. multiple rooms or a home), and also allow us to reduce carrier transmit power levels such that battery-powered mobile devices like smartphones can take the role of the central station.

Frequency-hopping in Morpho: The main issue to consider when extending our architecture to frequency-hopping spread spectrum radios is the fact that passive radios are not frequency selective. Frequency hopping can be enabled on the passive transmitter side (i.e. backscatter), by leveraging frequency shifting and recent advances in single-sideband backscattering [20]. These methods only add a small amount of complexity and power to our design. One area that needs more research to complete this design is the question of how to endow the passive receiver with similar frequency hopping capabilities.

One way to circumvent this issue is to use a dedicated channel for passive communication, and allow active communication to proceed with frequency hopping. The advantage of such decoupling is that we can use passive components like a SAW filter tuned to the specific channel before the passive receiver to make it frequency selective [17]. But this method restricts us to a single channel for passive communication which might be scalable to large networks.

10 RELATED WORK

We briefly review relevant related work that we have not highlighted in previous sections.

Multi-radio wireless networks: There has been much work on the general idea of multi-radio wireless radios. This work has explored many combinations including Bluetooth + WiFi [1, 5, 37], WiFi + LTE [7, 10, 26], and WiFi + 60 GHz [47]. This work has leveraged multi-radio combinations for energy-efficiency [1, 15, 21, 24, 25], traffic management [14], mobility management [38], and routing management [2, 11].

While there is similarity between these efforts and ours at a high level, the crucial difference is that we are designing an multi-radio system that operates at ultra-low power regimes between 1 μW to 1 mW, and can switch at microsecond granularity to react to highly dynamic channels. This is a completely different design space and necessitates rethinking all layers of the stack.

Active-Passive radios: There has been some recent work that explores integration of active and passive components, albeit in restricted ways. In terms of receiver side integration, recent work on wakeup radios integrate passive envelope detectors with active receivers to enable extremely low power remote wakeup [44, 45]. In terms of transmitter side integration, a recent short paper looks at reusing hardware elements between 10 Mbps BPSK Backscatter and 1Mbps Bluetooth [43]. A couple of approaches have explored integration at higher layers of the stack as well. One is BLISP [19], which we have previously discussed. Another is Braidio [17], which leverages active and passive components for power offload by shifting carrier generation between end-points. Our work shows how such active-passive radio components can be leveraged in every aspect of communication including data transfer, measurement, and control messages.

Backscatter communication: There has been significant activity in backscatter communication in recent years. A significant fraction of this work has focused on repurposing ambient carriers such as Bluetooth [12, 20, 57], WiFi [8, 23, 36, 54–57], Zigbee [20, 57], FM [51], and LoRA [48] to enable backscatter communication. Recent work has also shown that it is possible to use backscatter for applications like low-power HD video streaming [33]. However, the issue of how to deal with the inherent flakiness of passive radios under channel dynamics has received very little attention. Morpho bridges this gap.

11 CONCLUSION

In conclusion, we present a new architecture for low-power radios that leverages passive and active components in a tightly intertwined manner to improve performance. In contrast to duty-cycling based radios that aim to maximize sleep times to save power, polymorphic radios leverage passive modes to save power. This is a new paradigm that is particularly useful for low-power radios that are used in streaming mode to transmit data from or to wearable, IoT, and mobile devices. We instantiate our ideas in a full hardware-software stack that we call Morpho, and show that we can get up to an order of magnitude improvements in energy-efficiency while still being robust to channel fluctuations.

Acknowledgements

We thank Arun Venkataramani, our shepherd Sachin Katti, and the anonymous reviewers for their helpful feedback on the paper. This research was partially funded by NSF award #1719386, Intel Corp award #34627511 and ARL Cooperative Agreement W911NF-17-2-0196.

REFERENCES

[1] Y. Agarwal, T. Pering, R. Want, and R. Gupta. Switchr: Reducing system power consumption in a multi-client, multi-radio environment. In *Wearable Computers, 2008. ISWC 2008. 12th IEEE International Symposium on*, pages 99–102. IEEE, 2008.

[2] M. Alicherry, R. Bhatia, and L. E. Li. Joint channel assignment and routing for throughput optimization in multi-radio wireless mesh networks. In *Proceedings of the 11th annual international conference on Mobile computing and networking*, pages 58–72. ACM, 2005.

[3] Analog Devices. *0 Hz to 4.5 GHz, 40 dB Off Isolation at 1 GHz, 17 dBm P1dB at 1 GHz SPST Switches.* Rev. D.

[4] Analog Devices. *Wideband 4 GHz, 43 dB Isolation at 1 GHz, CMOS 1.65 V to 2.75 V, 2:1 Mux/SPDT.* Rev. E.

[5] G. Ananthanarayanan and I. Stoica. Blue-fi: enhancing wi-fi performance using bluetooth signals. In *Proceedings of the 7th international conference on Mobile systems, applications, and services*, pages 249–262. ACM, 2009.

[6] Avago Technologies. *Surface Mount Zero Bias Schottky Detector Diodes*, 5 2009.

[7] M. Bennis, M. Simsek, A. Czylwik, W. Saad, S. Valentin, and M. Debbah. When cellular meets wifi in wireless small cell networks. *IEEE communications magazine*, 51(6):44–50, 2013.

[8] D. Bharadia, K. R. Joshi, M. Kotaru, and S. Katti. Backfi: High throughput wifi backscatter. *ACM SIGCOMM Computer Communication Review*, 45(4):283–296, 2015.

[9] Bosch. *BMI160: Ultra Low Power Inertial Measurement Unit.*

[10] S. Deng, R. Netravali, A. Sivaraman, and H. Balakrishnan. Wifi, lte, or both?: Measuring multi-homed wireless internet performance. In *Proceedings of the 2014 Conference on Internet Measurement Conference*, pages 181–194. ACM, 2014.

[11] R. Draves, J. Padhye, and B. Zill. Routing in multi-radio, multi-hop wireless mesh networks. In *Proceedings of the 10th annual international conference on Mobile computing and networking*, pages 114–128. ACM, 2004.

[12] J. F. Ensworth and M. S. Reynolds. Every smart phone is a backscatter reader: Modulated backscatter compatibility with bluetooth 4.0 low energy (ble) devices. In *RFID (RFID), 2015 IEEE International Conference on*, pages 78–85. IEEE, 2015.

[13] Ettus Research. *USRP X300: High performance, Scalable, Software Designed Radio (SDR).*

[14] S. Ferlin, T. Dreibholz, and Ö. Alay. Multi-path transport over heterogeneous wireless networks: Does it really pay off? In *Global Communications Conference (GLOBECOM), 2014 IEEE*, pages 4807–4813. IEEE, 2014.

[15] J. Gummeson, D. Ganesan, M. D. Corner, and P. Shenoy. An adaptive link layer for heterogeneous multi-radio mobile sensor networks. *IEEE Journal on Selected Areas in Communications*, 28(7), 2010.

[16] Himax. *HM01B0: Ultra Low Power Image Sensor.*

[17] P. Hu, P. Zhang, M. Rostami, and D. Ganesan. Braidio: An integrated active-passive radio for mobile devices with asymmetric energy budgets. In *Proceedings of the 2016 conference on ACM SIGCOMM 2016 Conference*, pages 384–397. ACM, 2016.

[18] Impinj. *Impinj XArray RAIN RFID Gateway.*

[19] I. in't Veen, Q. Liu, P. Pawelczak, A. Parks, and J. R. Smith. Blisp: Enhancing backscatter radio with active radio for computational rfids. In *RFID (RFID), 2016 IEEE International Conference on*, pages 1–4. IEEE, 2016.

[20] V. Iyer, V. Talla, B. Kellogg, S. Gollakota, and J. Smith. Inter-technology backscatter: Towards internet connectivity for implanted devices. In *Proceedings of the 2016 conference on ACM SIGCOMM 2016 Conference*, pages 356–369. ACM, 2016.

[21] T. Jin, G. Noubir, and B. Sheng. Wizi-cloud: Application-transparent dual zigbee-wifi radios for low power internet access. In *INFOCOM, 2011 Proceedings IEEE*, pages 1593–1601. IEEE, 2011.

[22] S. Kamath and J. Lindh. Measuring bluetooth low energy power consumption. *Texas instruments application note AN092, Dallas*, 2010.

[23] B. Kellogg, V. Talla, S. Gollakota, and J. R. Smith. Passive wi-fi: Bringing low power to wi-fi transmissions. In *NSDI*, volume 16, pages 151–164, 2016.

[24] B. Kusy, C. Richter, W. Hu, M. Afanasyev, R. Jurdak, M. Brünig, D. Abbott, C. Huynh, and D. Ostry. Radio diversity for reliable communication in wsns. In *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on*, pages 270–281. IEEE, 2011.

[25] D. Lymberopoulos, N. B. Priyantha, M. Goraczko, and F. Zhao. Towards energy efficient design of multi-radio platforms for wireless sensor networks. In *Proceedings of the 7th international conference on Information processing in sensor networks*, pages 257–268. IEEE Computer Society, 2008.

[26] R. Mahindra, H. Viswanathan, K. Sundaresan, M. Y. Arslan, and S. Rangarajan. A practical traffic management system for integrated lte-wifi networks. In *Proceedings of the 20th annual international conference on Mobile computing and networking*, pages 189–200. ACM, 2014.

[27] Maxim Integrated. *Ultra-Low Power, Single-Channel Integrated Biopotential (ECG, R to R Detection) AFE.*

[28] A. Mayberry, P. Hu, B. Marlin, C. Salthouse, and D. Ganesan. ishadow: design of a wearable, real-time mobile gaze tracker. In *Proceedings of the 12th annual international conference on Mobile systems, applications, and services*, pages 82–94. ACM, 2014.

[29] C. A. P. Mercier, Patrick P. *Ultra-Low-Power Short-Range Radios.* Springer, 2015.

[30] MicroSemi. *IGLOO nano Low Power Flash FPGAs*, 9 2015. Rev. 19.

[31] Mini-Circuits. *Frequency Mixer, Level 4 (LO Power +4 dBm) 10 to 1000 MHz.* Rev. E.

[32] Mini-Circuits. *Power Splitter/Combiner 2 Way-0Âř 50Îľ 5 to 1000 MHz.* Rev. F.

[33] S. Naderiparizi, M. Hessar, V. Talla, S. Gollakota, and J. R. Smith. Low-power HD video streaming. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*, 2018.

[34] Nordic Semiconductor. *nRF52840: Ultra-low power 2.4GHz wireless system on chip (SoC).*

[35] NXP Semiconductors. *NPN wideband silicon RF transistor*, 3 2014. Rev. 2.

[36] A. N. Parks, A. Liu, S. Gollakota, and J. R. Smith. Turbocharging ambient backscatter communication. *ACM SIGCOMM Computer Communication Review*, 44(4):619–630, 2015.

[37] T. Pering, Y. Agarwal, R. Gupta, and R. Want. Coolspots: reducing the power consumption of wireless mobile devices with multiple radio interfaces. In *Proceedings of the 4th international conference on Mobile systems, applications and services*, pages 220–232. ACM, 2006.

[38] C. Pluntke, L. Eggert, and N. Kiukkonen. Saving mobile device energy with multipath tcp. In *Proceedings of the sixth international workshop on MobiArch*, pages 1–6. ACM, 2011.

[39] D. M. Pozar. *Microwave engineering.* John Wiley & Sons, 2009.

[40] PulseLarsen Antennas. *Penta Band Stubby Antenna*, 2 2010. Rev. 1.

[41] I. Recommendation. Vocabulary for performance and quality of service, 2006.

[42] J. M. Rehg, S. A. Murphy, and S. Kumar. *Mobile Health: Sensors, Analytic Methods, and Applications.* Springer, 2017.

[43] M. S. Reynolds. A 2.4-ghz, hybrid 10-mb/s bpsk backscatter and 1-mb/s fsk bluetooth tx with hardware reuse. *IEEE Microwave and Wireless Components Letters*, 27(12):1155–1157, 2017.

[44] N. E. Roberts, K. Craig, A. Shrivastava, S. N. Wooters, Y. Shakhsheer, B. H. Calhoun, and D. D. Wentzloff. 26.8 a 236nw- 56.5 dbm-sensitivity

bluetooth low-energy wakeup receiver with energy harvesting in 65nm cmos. In *Solid-State Circuits Conference (ISSCC), 2016 IEEE International*, pages 450–451. IEEE, 2016.

[45] N. E. Roberts and D. D. Wentzloff. Ultra-low power wake-up radios. In *Ultra-Low-Power Short-Range Radios*, pages 137–162. Springer, 2015.

[46] A. P. Sample, D. J. Yeager, P. S. Powledge, A. V. Mamishev, and J. R. Smith. Design of an rfid-based battery-free programmable sensing platform. *IEEE transactions on instrumentation and measurement*, 57(11):2608–2615, 2008.

[47] S. Sur, I. Pefkianakis, X. Zhang, and K.-H. Kim. Wifi-assisted 60 ghz wireless networks. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*, pages 28–41. ACM, 2017.

[48] V. Talla, M. Hessar, B. Kellogg, A. Najafi, J. R. Smith, and S. Gollakota. Lora backscatter: Enabling the vision of ubiquitous connectivity. *arXiv preprint arXiv:1705.05953*, 2017.

[49] A. Varshney, O. Harms, C.-P. Penichet, C. Rohner, F. Hermans, and T. Voigt. Lorea: A backscatter architecture that achieves a long communication range. In *ACM SenSys 2017*. ACM Digital Library, 2017.

[50] Vesper. *VM1010: Wake-on-Sound Piezoelectric MEMS Microphone*.

[51] A. Wang, V. Iyer, V. Talla, J. R. Smith, and S. Gollakota. Fm backscatter: Enabling connected cities and smart fabrics. In *NSDI*, pages 243–258, 2017.

[52] P. H. P. Wang, H. Jiang, L. Gao, P. Sen, Y. H. Kim, G. M. Rebeiz, P. P. Mercier, and D. A. Hall. A 400 mhz 4.5 nw -63.8 dbm sensitivity wake-up receiver employing an active pseudo-balun envelope detector. In *ESSCIRC 2017 - 43rd IEEE European Solid State Circuits Conference*, pages 35–38, Sept 2017.

[53] R. Want. Rfid explained: A primer on radio frequency identification technologies. *Synthesis Lectures on Mobile and Pervasive Computing*, 1(1):1–94, 2006.

[54] G. Yang, Y.-C. Liang, R. Zhang, and Y. Pei. Modulation in the air: Backscatter communication over ambient ofdm carrier. *arXiv preprint arXiv:1704.02245*, 2017.

[55] P. Zhang, D. Bharadia, K. Joshi, and S. Katti. Enabling backscatter communication among commodity wifi radios. In *Proceedings of the 2016 conference on ACM SIGCOMM 2016 Conference*, pages 611–612. ACM, 2016.

[56] P. Zhang, D. Bharadia, K. R. Joshi, and S. Katti. Hitchhike: Practical backscatter using commodity wifi. In *SenSys*, pages 259–271, 2016.

[57] P. Zhang, C. Josephson, D. Bharadia, and S. Katti. Freerider: Backscatter communication using commodity radios. In *Proceedings of the 13th International Conference on Emerging Networking EXperiments and Technologies*, CoNEXT '17, 2017.