

Sensor Assisted Wireless Communications

(Invited Paper)

Naveen Santhapuri* Justin Manweiler* Souvik Sen*
Xuan Bao* Romit Roy Choudhury* Srihari Nelakuditi†
*Duke University †University of South Carolina

Abstract

The nature of human mobility demands that mobile devices become agile to diverse operating environments. Coping with such diversity requires the device to assess its environment, and trigger appropriate responses to each of them. While existing communication subsystems rely on in-band wireless signals for context-assessment and response, we explore a lateral approach of using out-of-band sensor information. We propose a relatively novel framework that synthesizes in-band and out-of-band information, facilitating more informed communication decisions. We believe that further research in this direction could enable a new kind of device agility, deficient in today's communication systems. Since such a framework is located at the boundaries of mobile sensing and wireless communications, we call it sensor assisted wireless communications.

1. Introduction

In everyday life, mobile devices transition through a variety of environments. These transitions arise not only from changes in the physical environment (due to mobility), but also due to fluctuations within and across different wireless technologies. For instance, a mobile phone carried by an office-goer may transition from a stationary state (at home) to a walking state (on the way to the subway station) to a highly mobile state (within the train). The wireless background changes as well, subjecting the phone through WiFi (at home), 3G (outdoors), and then a period of complete disconnection, except perhaps when the train stops at the stations. Unsurprisingly, these environmental changes strain the communication subsystems in these devices. Wireless networking protocols must constantly discern the “context” of communication, and adapt with a new agility.

Context-aware device agility remains an elusive research challenge because the contexts are often difficult to discern based on wireless signal observations alone. A microwave oven at a Starbucks may induce the same kind of performance degradation as observed in a classroom with congested WiFi users. While the ideal response to the microwave would be to switch to a different frequency channel, less-frequent transmission attempts is most suitable for alleviating congestion. Clearly, such agility is difficult to attain without proper context assessment. Today's mobile devices optimize for the common case, and sacrifice performance where the context is atypical. This paper explores new opportunities for context-discrimination, thereby leading to improved device agility for truly pervasive communication.

Our main idea is simple. By using mobile phone sensors as an *out-of-band* context-assessment tool, we show that a certain kind of device agility may be achieved. For example, the phone's accelerometer measurements could identify the moving subway train, and switch off the WiFi/GSM subsystems to save energy. When the phone stops at each station, or when the user gets off the train, the accelerometer readings can pick the cue and switch on wireless access. As a generalization, the growing number of sensors on mobile devices presents an *out-of-band* opportunity to discern the communication context. While these contexts have been abundantly used in mobile computing applications [5], there is limited work that connects them to MAC/PHY layer functions [9]. This project attempts to make (and strengthen) this connection. Although an early work, our key contributions can be summarized as follows.

(1) *We identify the opportunity of utilizing out-of-band sensor information to optimize wireless*

communication systems (Figure 1). Of particular interest are the cases where the contextual information is implicitly present in the system.

(2) We propose a sensor-assisted wireless networking framework, and instantiate it with case-studies.

(3) We discuss our ongoing research on the generalization of these ideas. Our thought experiments show that this relatively new research space could be a lateral approach to augment existing research on wireless networking.

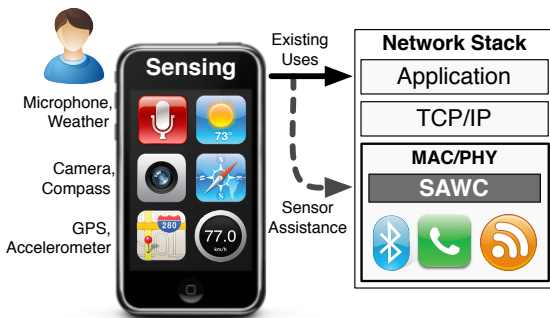


Figure 1. The Sensor Assisted Wireless Communications (SAWC) framework.

2. Motivation and Opportunity

This section zooms into the scope of *sensor assisted wireless communications*, and asks a set of natural questions. The intent is to characterize the need for sensor-assistance, gain an understanding of the opportunities, and therefrom, carve out the utility of the system. Three example applications are discussed as a verification of the general ideas in specific, real-world settings. Two of these applications are quantified through preliminary measurements and implementation.

We begin by asking, **why is in-band information inadequate for wireless protocol agility?** Our argument is two-fold. (1) Environmental factors and user mobility perturb wireless communication in diverse and unpredictable ways. When all these factors are further combined with interferences from nearby transmissions, the net observable result is purely stochastic. Diagnosing the causes of perturbations from the wireless signals is difficult. Moreover, the signal-level information exported by commodity wireless hardware is coarse-grained, making it

harder to segregate the causes. Nevertheless, the protocol's ideal reaction to the perturbation can vary depending on the cause of perturbation. An in-band approach to may not always be optimal.

A simple example of this can be seen in wireless rate adaptation. It has been well known that the transmission rate of a packet must be reduced only upon channel fading; the rate could remain the same (or even increase) if the packet loss was due to collision. Yet, today's wireless protocols are unable to reliably discriminate between fading and collision, and blindly reduce the rate upon packet losses. This is because *in-band* discrimination has been a difficult problem [10], [12]. As we elaborate later, using out-of-band information may facilitate failure diagnosis in this case.

(2) Our second argument pertains to the overhead of in-band approaches. Diagnosing the cause of signal perturbations may require bandwidth and energy investments, cutting back on the system's throughput and/or battery life. When the microwave oven is turned on in the vicinity, the WiFi interface may need to probe multiple channels before switching to the best one. Worse, the interface may need to continue probing (like polling mechanisms) to determine when it must switch back to the original channel. In-band probing will consume channel time, in addition to forcing frequent disassociations/reassociations to access points. Identifying the microwave through an out-of-band sound sensor could be a more effective method of channel switching (much like *interrupt* based operations). Since these sound sensors may anyway be active for a variety of other applications [1], [8], the cost of using them may get amortized.

Along this thread of argument, the natural follow-up question is **why would out-of-band techniques be any better for discriminating the context?** If the microwave's presence cannot be discerned through RF signals, it may not be discernible through sound either, i.e., other noises may drown the microwave-specific "hum". While that is true, we first clarify that in-band and out-of-band techniques are complementary, and could be used to balance out mutual deficiencies. Second, we argue that contextual events may have multiple fingerprints scattered over multiple sensing dimensions. Moving in a subway train may be identifiable through accelerometers and

sound sensors (trains have a characteristic acoustic signature easy to identify). Since mobile devices are equipped with a growing number of sensing dimensions, the context may be discernible over at least one of these dimensions. This diversity is likely to improve context-discrimination over single-dimensional, in-band techniques.

The next question then is **what is the space of opportunity for out-of-band techniques? What are example applications?** Figure 2 shows one possibility of classifying the broad topic of context-awareness from the perspective of wireless communications. The rows characterize the source of contextual information. When the contextual source belongs to the same sensing dimension, it is termed *in-band*, while information derived from other modes are *out-of-band*. Of course, one may argue that GPS location is an out-of-band information and has been abundantly used in improving communications [9], [5]. In light of this, we find that information can also be classified into those that are *explicitly* produced for consumers, and others that are *implicitly* present in the environment. For example, RTS/CTS packets in 802.11 are meant to *explicitly* alert nearby devices of an imminent transmission; GPS also offers *explicit* location information to enable context-awareness. The microwave “hum”, on the other hand, is not meant for a communication-related operation, and thereby an *implicit* source of information. A device’s secure identifier, extracted from inherent clock skews and frequency offsets [3], is again *implicit* for the same purposes. When viewed in this manner, we find that existing work has not adequately explored the class of *out-of-band, implicit* techniques. Yet, the number of sensing dimensions, as well as the modes of communications, continue to proliferate on mobile devices. This paper identifies this opportunity of information synthesis, and suggests a holistic approach to improve wireless network performance.

We now sample these higher level ideas through the following applications.

- Microwave-aware channel switching using out-of-band sound sensing.
- MAC layer bit rate adaptation using mobility and location classification.
- Augmenting end user’s communication experience using out-of-band activity recognition.

The above applications belong to the \langle *out-of-band, implicit* \rangle category, and are chosen to reflect

	Implicit	Explicit
In band	RadioTelepathy [Mobicom 08] Paradis [Mobicom 08]	AP Beacons, RTS/CTS BreadCrumbs [Mobicom 08]
Out of band	<i>MWave Interference Detection</i> <i>Accelerometer-based Rate Control</i>	GPS-based Rate Control [ICNP 08] Blue-Fi [Mobisys 09]

Figure 2. Classification of information sources from the perspective of mobile communications.

potential types of sensor assistance. We discuss them next and present preliminary validations.

3. Applications

3.1. Microwave-Aware Channel Switching

Opportunity: A running microwave oven can substantially degrade the performance of 2.4GHz 802.11 networks because they interfere with channels 6 through 11. However, given that most AP deployments select non-overlapping channels within the spatial vicinity, it may be beneficial for a device to switch to an AP that uses a non-overlapping channel (say, 1). Even though the link to the new AP could have a lower signal strength (RSSI), the switch may still be worthwhile. Once the microwave turns off, the device could switch back to channel 6 or 11.

Context-change Detection: As mentioned earlier, detecting the microwave’s signature, in-band, may not be simple. Ongoing 802.11 packets from hidden terminals, variable packet lengths, channel noise, and non-802.11 interferences from Bluetooth/cordless-phones systems are likely to complicate diagnosis. These stochastic effects make context-assessment inherently ambiguous, resulting in false alarms. Even if the microwave is somehow detected, the subsequent difficulty arises in knowing the optimal time to switch back to the original channel. The obvious technique would be to periodically probe for improved conditions on the original, microwave-affected channel. However, probing entails disassociation from the current AP and re-association to the original AP. This is a costly operation and can be prohibitive if triggered frequently. If infrequently triggered, the device may not realize that the microwave has stopped, and will unnecessarily remain on the sub-optimal channel for the long probing period.

A sound sensor can provide out-of-band information regarding the presence of a nearby microwave in use. The characteristic background “hum” or the familiar microwave beeps can provide an acoustic signature. Figure 3 shows this signature on the frequency domain. The microwave sounds were recorded using a Nokia N95 phone (at distances of 0, 1, 3, 5, and 15 meters) and subjected through simple analysis. Across all tests, we found reliable detection accuracy, with false negatives/false positives at 1.5% and 4.6%, respectively. We believe this could enable the desired environment-aware agility in future mobile devices.




Figure 3. Microwave oven’s acoustic signature on the frequency domain

Throughput Improvement: Given that sound-based channel switching is feasible, we tested potential performance gains from avoiding microwave interference. Using `iperf`, we tested TCP throughput across the commonly-used channels 1, 6, and 11. We used three microwaves of differing age and vendor, operating at center frequency of 2.45 GHz, between 802.11b/g channels 9 and 10. The microwave was placed at varying distances from the TCP receiver, which was then switched to different channels. To maximize consistency across channels, we conducted all experiments at night, minimizing network contention. Figure 4(a, b, c) show the impact of microwaves (MW) on throughput, when the laptops were switched to channels 1, 6, and 11, respectively. For a switch from channel 11 to 1, we find an average improvement of 83% and 87% under interference from MW1 and MW2, respectively. A switch from channel 6 to 1 offers 87% and 75% under MW1 and MW2. Results from MW3 are similar and omitted for visual clarity. We believe these gains are substantial, especially considering that only a simple channel change is sufficient to realize them.

3.2. Sensor-aware Rate Control

Opportunity: Wireless rate adaptation is a challenging problem. Among other reasons, the difficulty in estimating channel fluctuations is the most prominent one. The problem is exacerbated because the nature of the fluctuations vary under different mobility regimes. Research has shown that the effects of path loss dominate in a vehicular network scenario, while multipath effects are strong in indoor low-mobility scenarios [11]. One-size-fits-all rate control protocols are difficult to design, and hence, any given protocol is best suited for a subset of the mobility regimes. Specifically, SNR-based rate control algorithms are shown to perform well in urban outdoor environments [4]. In contrast, packet error-based schemes like `SampleRate` are suitable for static indoor environments [12]. A simple way to distinguish indoor/outdoor locations is to obtain the light sensor readings on the phone and classify them based on time of the day. We also suggest the possibility of identifying the mobility regime of the user (static, walking, vehicular) through on-phone accelerometers. This information could be valuable in appropriately multiplexing between rate control algorithms (when feasible), or simply optimizing the given protocol for that condition.

Context-change Detection: To verify the possibility of detecting the user’s mobility regime, we drove a sedan in a 25-mile loop at an average speed of 28 MPH. The route consisted of a representative mix of both well and poorly-paved roads. Three Nokia N95s continually recored accelerometer readings. Fig. 5 shows the standard deviation of acceleration across a portion of one trace. Clear patterns of peaks and valleys emerge, reflecting movement and stillness, respectively. We also tested the accelerometer signature when a user walked with the phone in her pocket. A small portion of the trace shows distinct rhythmic patterns (Fig. 5). Such pattern [8] has been attempted before but we only show these results to illustrate the opportunity. We believe that this can provide the out-of-band information necessary for informed rate control decisions.

3.3. Activity Assisted Communications

Taking a broad view of mobile communication, we consider the human element in maintaining

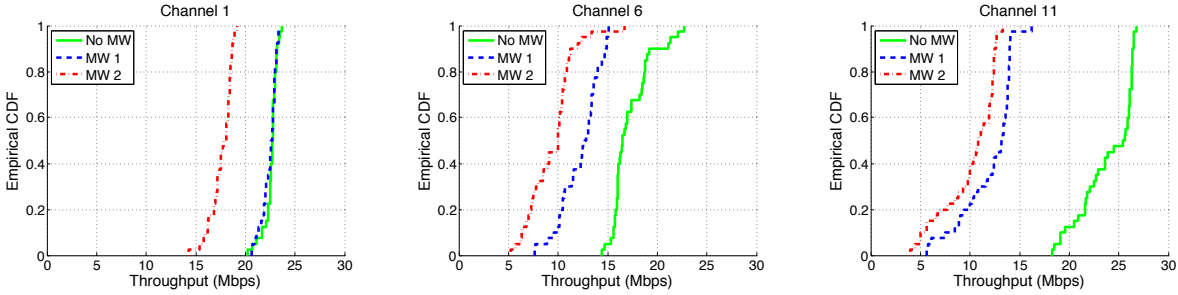


Figure 4. Throughput comparison across 802.11b/g channels 1, 6, 11. The impact of microwave interference is substantially less on channel 1.




Figure 5. An acceleration profile distinguishing vehicular motion from stopped periods; Clearly distinguishable walking patterns

a wireless channel. Ultimately, end-user experience is the metric-of-interest. Pervasive activity recognition might improve the convenience of wireless communications. Suppose a phone call is inopportunistly received while driving. Instead of ringing, the agile device might instead confirm with the caller first, “The person you’ve dialed is driving. Do you wish to continue?” Moreover, if the caller declines, the question has provided a helpful cue as to how long the caller should wait before trying later. We implemented this simple application on a Google Nexus One Phone. If the callee is in a moving car and the call was not answered, our preliminary application senses the mobility of the user and sends a text message to the caller. Clearly, many situations can benefit from out-of-band contexts (e.g., watching a movie, attending a seminar, jogging). Office phone systems provide a rudimentary attempt, but require human participation to setup (e.g., “in a meeting,” “on vacation”). Advanced activity sensing combined with pervasive localization (e.g., [8] and [1]) can provide a deeper context awareness that would be difficult to extract in-band. Sharing this context with the callee helps to optimize the *experience* of the communication.

Of course, there are privacy implications for automatically sharing contexts between users. How-

ever, as has been already shown in a number of existing systems (e.g., Facebook, Google Latitude), straightforward configurability can balance utility with risks. We believe that the integration of human activity recognition and direct feedback can provide useful contexts, improving the end-user experience.

4. Related Work

Several works use sensors to characterize the ambience [1], [7]. We consider sensor assistance to infer the RF environment, allowing optimized wireless performance. In-band schemes, such as channel hopping [6], have been proposed to avoid non-compliant interference (e.g., from cordless phones, microwaves). Our work is complementary. Inferred context from out-of-band channels can enhance known remedies. Using sensor information for improving wireless performance is relatively unexplored. Context-aware rate control uses GPS location and history to infer pathloss, thereby adjusting wireless bitrate [11]. This is an explicit out-of-band inference mechanism. We are not aware of any work utilizing out-of-band implicit channels for improving wireless *performance*. However, [9] has considered the use of context-sensing in optimizing the energy use of a mobile device, switching between WiFi and GSM connectivity as appropriate.

5. Limitations and Discussion

Out-of-band context may not always be discernible and accurate, and may incur additional cost and latency. This section discusses some guidelines to employ sensor assisted communication.

Out-of-band information should supplement, not supplant, in-band information. While diverse

sensors for sound, light, speed, etc. provide out-of-band information in multiple dimensions, it may still not be adequate to discern context in some environments. For example, white space networking [2] permits secondary users to reuse a spectrum, provided they do not interfere with primary users, say microphones. We can facilitate such an opportunistic spectral reuse if the presence of microphones can be detected through sound sensors. However, inaudible microphones can still interfere with secondary users. Moreover, when out-of-band information is available, it should be coupled with in-band information to determine the appropriate course of action. In other words, out-of-band information should be treated as a helpful hint to address a problem rather than a complete solution.

Out-of-band information should provide proper and timely context. One could argue that the information obtained from out-of-band channels must have high fidelity. Otherwise, it could cause the protocol operations to be tuned for an incorrect context. Discerning the context in a timely manner is also necessary to respond to frequent contextual transitions. Our case studies have shown that we can identify an active microwave and a moving car with reasonable accuracy and latency. Note that even when the out-of-band information is not quite precise, it may still provide useful hints for protocol adaptation.

Overhead of out-of-band information should ideally be minimal. The ability to discern the context of communication with out-of-band information does not warrant its use in all situations. Even if the information is helpful, the cost of obtaining it in the out-of-band channel must be less than the additional cost of obtaining the same information in-band. But in many instances, the sensory information may be available at no additional cost since the sensors are typically always on to serve several other applications. Only when the desired sensory information is not immediately available, we need to assess the pros and cons of activating sensors. We believe that this cost can often be amortized over many other applications that benefit from context awareness.

6. Conclusion

Mobile devices must continually cope with challenging and diverse operating environments. With

their integration of sensing, computation, and wireless connectivity, modern mobile devices are uniquely positioned to characterize their surroundings. We believe that sensing can provide a necessary contextual awareness, allowing these devices to become truly agile. Ultimately, wireless systems may be made more robust to their environments, yielding an enhanced end-user experience.

References

- [1] M. Azizyan, I. Constandache, and R. R. Choudhury. Surroundsense: Mobile phone localization via ambient fingerprinting. In *MOBICOM*, 2009.
- [2] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh. White space networking with wi-fi like connectivity. In *SIGCOMM*, 2009.
- [3] V. Brik, S. Banerjee, M. Gruteser, and S. Oh. Wireless device identification with radiometric signatures. In *Proc. ACM Mobicom*, 2008.
- [4] J. Camp and E. Knightly. Modulation Rate Adaptation in Urban and Vehicular Environments: Cross-layer Implementation and Experimental Evaluation. In *MobiCom*, 2008.
- [5] G. Chen and D. Kotz. A survey of context-aware mobile computing research. *Dartmouth Computer Science Technical Report TR2000-381*, 2000.
- [6] R. Gummadi, D. Wetherall, B. Greenstein, and S. Seshan. Understanding and mitigating the impact of rf interference on 802.11 networks. In *Proc. ACM Sigcomm*, 2007.
- [7] H. Lu, W. Pan, N. D. Lane, T. Choudhury, and A. T. Campbell. SoundSense: Scalable Sound Sensing for People-Centric Sensing Applications on Mobile Phones. In *Mobisys*, 2009.
- [8] E. Miluzzo, N. Lane, K. Fodor, R. Peterson, H. Lu, M. Musolesi, S. Eisenman, X. Zheng, and A. Campbell. Sensing meets mobile social networks: The design, implementation and evaluation of the cenceme application. In *SENSYS*, 2008.
- [9] A. Rahmati and L. Zhong. Context-for-wireless: Context-sensitive energy-efficient wireless data transfer. In *MOBISYS*, 2007.
- [10] S. Rayanchu, A. Mishra, D. Agrawal, S. Saha, and S. Banerjee. Diagnosing Wireless Packet Losses in 802.11: Separating Collision from Weak Signal. In *INFOCOM*, 2008.
- [11] P. Shankar, T. Nadeem, J. Rosca, and L. Iftode. CARS: Context Aware Rate Selection for Vehicular Networks. In *ICNP*, 2008.
- [12] M. Vutukuru, H. Balakrishnan, and K. Jamieson. Cross-layer wireless bit rate adaptation. In *ACM SIGCOMM*, 2009.