

Application-Layer Clock Synchronization for Wearables Using Skin Electric Potentials Induced by Powerline Radiation

Zhenyu Yan

School of Computer Science and Engineering
Nanyang Technological University
Singapore
zyan006@ntu.edu.sg

Rui Tan*

School of Computer Science and Engineering
Nanyang Technological University
Singapore
tanrui@ntu.edu.sg

Yang Li

Advanced Digital Sciences Center
Illinois at Singapore
Singapore
yang.li@adsc.com.sg

Jun Huang

Center for Energy Efficient Computing and Applications
Peking University
Haidian, Beijing, China
jun.huang@pku.edu.cn

ABSTRACT

Design of clock synchronization for networked nodes faces a fundamental trade-off between synchronization accuracy and universality for heterogeneous platforms, because a high synchronization accuracy generally requires platform-dependent hardware-level network packet timestamping. This paper presents *TouchSync*, a new indoor clock synchronization approach for wearables that achieves millisecond accuracy while preserving universality in that it uses standard system calls only, such as reading system clock, sampling sensors, and sending/receiving network messages. The design of *TouchSync* is driven by a key finding from our extensive measurements that the skin electric potentials (SEPs) induced by powerline radiation are salient, periodic, and synchronous on a same wearer and even across different wearers. *TouchSync* integrates the SEP signal into the universal principle of Network Time Protocol and solves an integer ambiguity problem by fusing the ambiguous results in multiple synchronization rounds to conclude an accurate clock offset between two synchronizing wearables. With our shared code, *TouchSync* can be readily integrated into any wearable applications. Extensive evaluation based on our Arduino and TinyOS implementations shows that *TouchSync*'s synchronization errors are below 3 and 7 milliseconds on the same wearer and between two wearers 10 kilometers apart, respectively.

CCS CONCEPTS

• **Networks** → **Time synchronization protocols**; • **Human-centered computing** → **Ubiquitous and mobile devices**;

*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

SenSys '17, November 6–8, 2017, Delft, The Netherlands

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-5459-2/17/11...\$15.00

<https://doi.org/10.1145/3131672.3131681>

KEYWORDS

Clock synchronization, wearables, skin electric potential

ACM Reference Format:

Zhenyu Yan, Yang Li, Rui Tan, and Jun Huang. 2017. Application-Layer Clock Synchronization for Wearables Using Skin Electric Potentials Induced by Powerline Radiation. In *Proceedings of 15th ACM Conference on Embedded Networked Sensor Systems (SenSys'17)*. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3131672.3131681>

1 INTRODUCTION

The annual worldwide shipments of consumer wearables (e.g., smart watches, wristbands, eyewears, clothing, etc) have grown by 29% in 2016 [11]. This rapid growth is expected to continue, projecting to 213 million units shipped in 2020 [11]. Along with the proliferation of consumer wearables, specialized domains such as clinical/home healthcare [1] and exercise/sport analysis [25] are also increasingly adopting smart wearable apparatuses. In the body-area networks formed by these wearables, a variety of system functions and applications depend on tight clock synchronization among the nodes. For instance, two earbuds of a wireless headphone need to be synchronized mutually and/or with a master device (e.g., a smartphone) to control the playback positions in their buffers to deliver audio synchronously [3]. Motion analysis [22] and muscle activity monitoring [24, 25] require sensory data from multiple tightly synchronized nodes.

While current wearable systems adopt customized, proprietary clock synchronization approaches [8], we envisage a wide spectrum of interoperable wearables that can synchronize with each other to enable more novel applications. For instance, in body sensor assisted multi-user gaming that may need to decide which participant performs an action or gesture first, tight clock synchronization among the body sensors and/or the handheld game consoles is needed. In the envisaged scheme, an application developer can readily synchronize any two communicating wearables using high-level and standard system calls provided by their operating systems (OSes), such as reading system clock, transmitting and receiving network messages. However, the design of clock synchronization approaches faces a fundamental trade-off between the synchronization accuracy and the universality for heterogeneous platforms.

This is because a high synchronization accuracy generally requires low-level timestamping for the synchronization packets, which may be unavailable on the hardware platforms or inaccessible to the application developer.

We illustrate this accuracy-universality trade-off using the Network Time Protocol (NTP) [26] and the Precision Time Protocol (PTP) [12]. NTP synchronizes a slave node and a master node by recording their clock values when a UDP synchronization packet is passed to and received from the sender's and receiver's OSes, respectively. Thus, NTP is universal in that it can be applied to any host that speaks UDP. However, as the application-layer timestamping cannot capture the details of the nondeterministic OS and network delays, NTP may yield significant synchronization errors up to hundreds of milliseconds (ms) in a highly asymmetric network. To solve this issue, PTP uses hardware-level timestamping provided by PTP-compatible Ethernet cards at the end hosts and all the switches on the network path to achieve microsecond (μ s) accuracy. However, the need of the special hardware inevitably negates its universality and restricts PTP's adoption to time-critical local-area networks only, e.g., those found in industrial systems.

In wireless networks, due to the more uncertain communication delays caused by media access control (MAC), NTP performs worse. Thus, similar to PTP, most existing clock synchronization approaches for wireless sensor networks (WSNs) (e.g., RBS [4], TPSN [7], and FTSP [23]) have resorted to MAC-layer timestamping provided by the nodes' radio chips to pursue synchronization accuracy. While FTSP has become the *de facto* standard in TinyOS-based WSNs, the need of the MAC-layer timestamping presents a significant barrier for its wide adoption to the broader Internet of Things (IoT) domain, where the IoT platforms use diverse radios and OSes, and in general they do not provide an interface for the MAC-layer timestamping.

In this paper, we aim at developing a new clock synchronization approach for wearables that establishes a desirable accuracy-universality trade-off point between the two extremes represented by NTP and PTP to well address the momentum of IoT platform heterogenization. In particular, we will stem from the sensor nature of wearables to explore ambient signals that can assist clock synchronization. Recent studies exploited external periodic signals such as powerline radiation [28] and Radio Data System (RDS) [19] to calibrate the clocks of WSN nodes. However, these approaches need non-trivial extra hardware to capture the external signals. Moreover, they focus on *clock calibration* that ensures different clocks advance at the same speed, rather than synchronizing the clocks to have the same value. But they inspire us to inquire (i) the existence of a periodic and synchronous signal that can be sensed by different wearables without adding non-trivial hardware to preserve universality and (ii) how to exploit the signal to synchronize the wearables without using hardware-level packet timestamping.

For the first inquiry, we conduct extensive measurements to explore such signals. Our measurements based on Adafruit's Flora [14], an Arduino-based wearable platform, show that by simply sampling an onboard analog-to-digital converter (ADC), a Flora can capture powerline electromagnetic radiation that oscillates at the power grid frequency (e.g., 50 Hz). When the Flora's ADC has a physical contact with the wearer's skin, the sampled skin electric potential (SEP) is a significantly amplified version of the powerline

radiation, because the human body acts as an effective antenna. Although the SEP's amplitude is dynamic due to the human body movements, its frequency is highly stable. Moreover, the SEPs on the same wearer and even different wearers in a same indoor environment exhibit desirable synchronism. The time displacement between the SEPs at different positions of a still wearer is generally less than 1 ms. These results suggest that SEP is a promising basis for synchronizing the wearables.

For the second inquiry, we integrate the periodic and synchronous SEP signal into the universal principle of NTP to deal with the major source of NTP's error, i.e., asymmetric communication delays. In the original NTP, the problem of estimating the offset between the slave's and master's clocks is a *real-domain* underdetermined problem that has infinitely many solutions. NTP chooses a solution by assuming symmetric communication delays, which does not hold in many scenarios, however. Assisted with the periodic SEP, the problem reduces to an *integer-domain* underdetermined problem that has a finite number of solutions. However, it is challenging to infer which solution is correct. In this paper, we show that, if the time displacement between the SEP signals captured by two synchronizing wearables is shorter than half of the power grid voltage cycle, the integer ambiguity can be resolved by jointly considering multiple synchronization rounds with dynamic and asymmetric communication delays. Thus, the clock offset between a pair of wearables can be estimated with ms accuracy due to the ms synchronism between their SEP signals.

Based on the above two key results, we design a novel clock synchronization approach for wearables, which we call *TouchSync*. It runs at the application layer in that the needed SEP sampling, the network message exchange and timestamping can be implemented using standard wearable OS calls. Thus, by introducing a rather simple skin contact, we can readily achieve ms synchronization accuracy across heterogeneous wearable platforms, without resorting to the hardware-level packet timestamping that is extremely difficult to be standardized. To simplify the adoption of TouchSync by application developers, we design and release a `touchsync.h` header file [31] that implements TouchSync's signal processing algorithms and the integer ambiguity solver. With this header, the implementations of TouchSync in Arduino and TinyOS need about 50 and 150 lines of code, respectively, which manage sensor sampling, synchronization message exchange and application-layer timestamping only. All these tasks are basics for Arduino and TinyOS developers. We also conduct extensive experiments in various indoor environments to show the pervasive availability of the SEP signals. On the same wearer, TouchSync mostly achieves sub-ms accuracy and the largest error is 2.9 ms. We also conduct a TouchSync-over-Internet proof-of-concept experiment that yields errors below 7 ms between two wearers 10 km apart.

The ADC-skin contact needed by TouchSync can be easily integrated into the wearable designs with near-zero cost. Although our experiments show that, in the absence of the contact, TouchSync can still work with graceful performance degradation, SEP is a new and valuable sensing modality for integration consideration, since it provides accurate timing and is indicative of other information about the wearer (e.g., body orientation, movements, and indoor location) as suggested by our measurements in this paper.

The rest of this paper is organized as follows. Section 2 reviews related work. Section 3 introduces background and our objective. Section 4 presents a measurement study. Section 5, Section 6, and Section 7 designs, implements, and evaluates TouchSync, respectively. Section 8 concludes.

2 RELATED WORK

Highly stable time sources are ill-suited for wearables. Chip-scale atomic clocks are too expensive (\$1,500 per unit [29]). GPS receivers are power-consuming and do not work in indoor environments. Recent studies exploited external signals available in indoor environments to synchronize or calibrate the clocks of distributed nodes. In [2], an AM radio receiver is designed to decode the global time information broadcast by timekeeping radio stations. In [19], a mote peripheral is designed to capture the periodic RDS signals of FM radios for clock calibration. In [28], a mote peripheral called syntonistor can receive the periodic electromagnetic radiation from powerlines to calibrate wireless sensors' clocks, where some clock synchronization approach is still needed for the initial synchronization. In [27, 30], voltage sensors plugged in to wall power outlets are used to secure clock synchronization against malicious network delays. In particular, the voltage cycle length fluctuations are exploited to implement a data-based clock synchronization approach [30]. In [20], such fluctuations extracted from the powerline radiation are used as *natural timestamps*. However, the error of the natural timestamps can be up to hundreds of ms. Moreover, the clock synchronization based on the natural timestamps needs to transmit a considerable amount of cycle length data and a compute-intensive matching process to decode the fluctuations to time information [20]. Thus, the natural timestamping approach is ill-suited for tight clock synchronization among resource-constrained wearables. In [18], a smartphone captures ultrasonic beacons from pre-deployed synchronized beacon nodes to synchronize its own clock. All the above approaches [2, 18–20, 27, 28, 30] need non-trivial customized hardware and infrastructures, which reduce their universality.

Two recent studies leverage built-in sensing modalities to capture external periodic signals for clock calibration. In [10], a 802.15.4 radio is used to capture the Wi-Fi beacons to calibrate motes' clocks. Although this approach does not require a peripheral, it uses the received signal strength indication (RSSI) register of the CC2420 radio chip, which makes it hardware specific and nonuniversal for IoT platforms that use diverse radios. In [21], motes use light sensors to capture the fluorescent light that flickers at a frequency twice of the power grid frequency to calibrate their clocks. Although light sensors are widely available, the required fluorescent lighting may not be available in natural lighting environment. In contrast, the powerline radiation that induces the SEP signal used by our approach pervades civil infrastructures.

The studies [10, 19, 21, 28] mentioned above, including the two [21, 28] that are power grid related, focus on *clock calibration* that involves no message exchanges among nodes. Though continuous clock calibration maintains the nodes synchronized once they are initially synchronized, the initial synchronization and the resynchronizations needed upon clock calibration faults are not addressed in these studies. Thus, these studies and ours are complementary, in that the principle of TouchSync can be used for

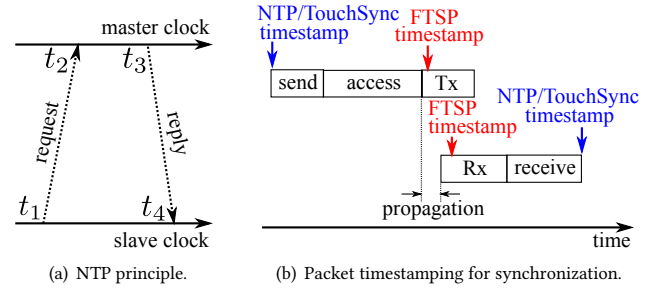


Figure 1: NTP principle and packet timestamping.

the initial synchronization of the systems adopting these clock calibration approaches [10, 19, 21, 28].

3 BACKGROUND AND OBJECTIVE

3.1 NTP Principle and Packet Timestamping

Many clock synchronization approaches adopt the principle of NTP, which is illustrated in Fig. 1(a). A *synchronization session* consists of the transmissions of a request packet and a reply packet. The t_1 and t_4 are the slave's clock values when the request and reply packets are transmitted and received by the slave node, respectively. The t_2 and t_3 are the master's clock values when the request and reply packets are received and transmitted by the master node, respectively. Thus, the round-trip time (RTT) is $RTT = (t_4 - t_1) - (t_3 - t_2)$. Based on a *symmetric link assumption* that assumes identical times for transmitting the two packets, the offset between the slave's and the master's clocks, denoted by δ_{NTP} , is estimated as $\delta_{NTP} = t_4 - (t_3 + \frac{RTT}{2})$. Then, this offset is used to adjust the slave's clock to achieve clock synchronization. Under the above principle, non-identical times for transmitting the two packets will result in an error in estimating the clock offset. The estimation error is half of the difference between the times for transmitting the two packets.

We use Fig. 1(b) and the terminology in [23] to explain how the timestamps (e.g., t_3 and t_4) are obtained in NTP and existing WSN clock synchronization approaches. The *send time* and the *receive time* are the times used by the OS to pass a packet between the synchronization program and the MAC layer at the sender and receiver, respectively. They depend on OS overhead. The *access time* is the time for the sender's MAC layer to wait for a prescribed time slot in time-division multiple access (TDMA) or a clear channel in carrier-sense multiple access with collision avoidance (CSMA/CA). It often bears the highest uncertainty and can be up to 500 ms [23]. The transmission (Tx) and reception (Rx) times are the physical layer processing delays at the sender and receiver, respectively. The *propagation time* equals the distance between the two nodes divided by the speed of light, which is generally below 1 μs .

As illustrated in Fig. 1(b), NTP timestamps the packet when the packet is passed to or received from the OS. Thus, the packet transmission time used by NTP is subjected to the uncertain OS overhead and MAC. Therefore, as measured in Section 4.2, NTP over a Bluetooth connection can yield nearly 200 ms clock offset estimation errors. To remove these uncertainties, FTSP uses MAC-layer access to obtain the times when the beginning of the packet

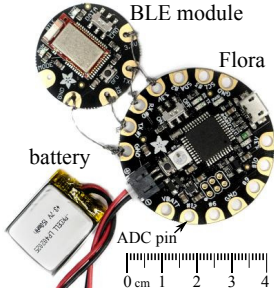


Figure 2: Flora.

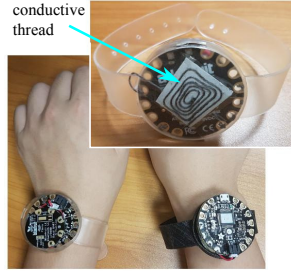


Figure 3: Prototypes supporting TouchSync.

is transmitted/received by the radio chip. As the propagation time is generally below $1\ \mu\text{s}$, FTSP simply estimates the clock offset as the difference between the two hardware-level timestamps. Thus, the two-way scheme in Fig. 1(a) becomes non-essential for FTSP.

3.2 Objective

NTP, though universal, gives unacceptably low accuracy. On the other extreme, existing WSN clock synchronization approaches, though achieving μs accuracy, may not be universally applicable to the diversified IoT platforms with different radio chips and OSes. In this paper, by introducing the readily available SEP signal, we aim at developing a new clock synchronization approach for wearables that (i) uses application-layer timestamping as NTP does to preserve universality and (ii) achieves ms accuracy that meets the requirements of a range of applications. For instance, music streaming requires a synchronism below 30 ms between two wireless earbuds [3]. For seismic sensing based motion analysis [22] and muscle activity monitoring [25] that generally adopt sampling rates up to 500 Hz, the ms synchronization accuracy can enable us to discriminate any readings sampled by different sensors at different time instants. To this end, we need to understand the properties of SEP, which is the subject of Section 4.

4 MEASUREMENT STUDY

In this section, we conduct measurements to gain insights for guiding the design of TouchSync.

4.1 Measurement Setup

Our measurement study uses two Adafruit Flora nodes [14] and a Raspberry Pi (RPI) 3 Model B single-board computer [5]. The Flora is an Arduino-based wearable platform that can be programmed using the Arduino IDE. Each Flora node, as shown in Fig. 2, consists of a main board with an ATmega32u4 MCU (8 MHz, 2.5 KB RAM), a Bluetooth Low Energy (BLE) 4.1 module, and a 150mAh lithium-ion polymer battery. The RPI has a built-in BLE 4.1 module and runs Ubuntu MATE 16.04 with BlueZ 5.37 as the BLE driver. We use Adafruit's nRF51 Arduino library [15] and BluefruitLE Python library [16] on the Floras and the RPI, respectively, to send and receive data over BLE in the UART mode through the `write()` and `read()` functions. The Floras and RPI operate as BLE peripheral (slave) and central (master), respectively. To obtain the ground truth clocks, in each experiment, we synchronize the Floras with the RPI

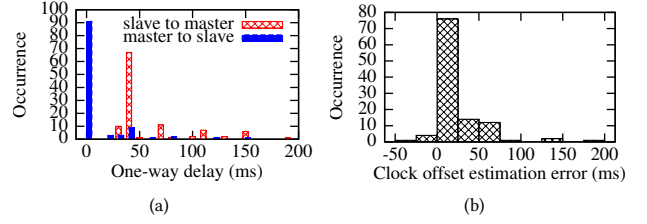


Figure 4: Performance of NTP over a BLE connection.

as follows. At the beginning of the experiment, we wire a general-purpose input/output (GPIO) pin of the RPI with a digital input pin of each Flora. Then, the RPI issues a rising edge through the GPIO pin and records its clock value t_{master} . Upon detecting the rising edge, a Flora records its clock value t_{slave} and sends it to the RPI. The RPI computes the ground-truth offset between the Flora's and the RPI's clocks as $\delta_{GT} = t_{\text{slave}} - t_{\text{master}}$. Then, we remove the wiring and conduct experiments.

4.2 Performance of NTP over BLE Connection

As our objective is to devise a new clock synchronization approach that uses application-layer timestamping as NTP does, this section measures the performance of NTP to provide a baseline. We implement the NTP principle described in Section 3.1 on the Flora setup. Fig. 4(a) shows the distributions of the one-way application-layer communication delays over 110 NTP sessions. The slave-to-master delays are mostly within $[40, 50]$ ms, with a median of 42 ms and a maximum of 376 ms (not shown in the figure). As specified by the BLE standard, the master device pulls data from a slave periodically. The period, called *connection interval*, is determined by the master. The slave needs to wait for a pull request to transmit a packet to the master. In BlueZ, the connection interval is set to 67.5 ms by default. As the arrival time of a packet from the slave's OS is uniformly distributed over the connection interval, the expected access time is $67.5/2 = 33.75$ ms. This is consistent with our measured median delay of 42 ms, which is about 8 ms longer because of other delays (e.g., send and receive times). The exceptionally long delays (e.g., 376 ms) observed in our measurements could be caused by transient wireless interference and OS delays. For the master-to-slave link, the delays are mostly within $[0, 10]$ ms, with a median of 8 ms and a maximum of 153 ms. A BLE slave can skip a number of pull requests, which is specified by the *slave latency* parameter, and sleep to save energy. Under BlueZ's default setting of zero for slave latency, the slave keeps awake and listening, yielding short master-to-slave delays.

The asymmetric slave-to-master and master-to-slave delays will cause significant errors in the NTP's clock offset estimation. At the end of each synchronization session, the RPI computes this error as $\delta_{NTP} - \delta_{GT}$, where δ_{NTP} and δ_{GT} are NTP's estimate and the ground-truth offset, respectively. Fig. 4(b) shows the distribution of the errors. We observe that 28% of the errors are larger than 25 ms. The largest error in the 110 NTP sessions is 183 ms. Such an error profile does not well meet the ms accuracy requirements of many applications [3, 22, 25]. Though it is possible to calibrate the average error to zero by using prior information (e.g., the settings of

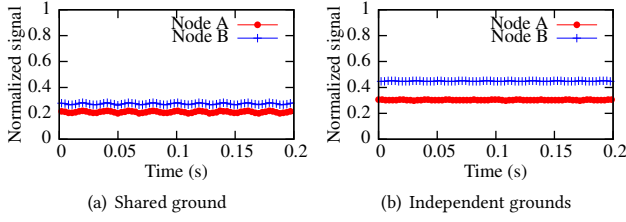


Figure 5: No human body contact.

the connection interval and slave latency), the calibration is tedious, nonuniversal, and incapable of reducing noise variance.

4.3 Skin Electric Potential (SEP)

In this set of measurements, we explore i) whether a human body is an effective antenna for receiving the powerline radiation and ii) whether the SEP signals induced by the radiation on the same wearer or different wearers are synchronous. The Flora's MCU has a 10-bit ADC that supports a sampling rate of up to 15 kHz. To facilitate experiments, we have made two Flora-based prototypes as shown in Fig. 3. We place the Flora into a 3D-printed insulating wristband and use a stainless thin conductive thread to create a connection between Flora's ADC pin and the wearer's skin. The Flora samples the ADC at 333 Hz continuously for two minutes and streams the timestamped raw data to the RPi for offline analysis. The sampling rate of 333 Hz is sufficient to capture the powerline radiation or SEP with a frequency of 50 Hz in our region. All samples are normalized using the reference voltage of the ADC. As only the ADC pin is connected to the researcher, the grounding of the Flora may affect the sampling result. To understand the impact of grounding, we conduct two sets of comparative experiments, where the two Floras have shared and independent grounds, respectively. In each experiment set, there are two scenarios: still and moving. For the moving scenario, the researcher keeps changing the body orientation, movement, and location. The experiments are conducted in a computer science laboratory with various appliances such as lights, computers, and printers.

4.3.1 Shared ground. We wire the ground pins of the two Floras, such that they have a shared ground. We conduct three tests.

First, Fig. 5(a) shows the signals captured by the two Floras when they have no physical contact with any human body. The signals have small fluctuations with a normalized peak-to-peak amplitude of 0.024. The signals fluctuate at a frequency of 50 Hz. This suggests that the Floras can pick up the powerline radiation. However, the signals are weak.

Second, a researcher touches the ADC pins of the two Floras with his two hands, respectively. Figs. 6(a) and 6(b) show the signals captured by the two nodes during the same time duration, when the researcher stands still and walks, respectively. Under the two scenarios (still and moving), the two nodes yield salient and almost identical signals. The peak-to-peak amplitudes in the two figures are around 0.4 and 0.8, which are 17 and 33 times larger than that of the signal shown in Fig. 5(a). This suggests that the human body can effectively receive the powerline radiation.

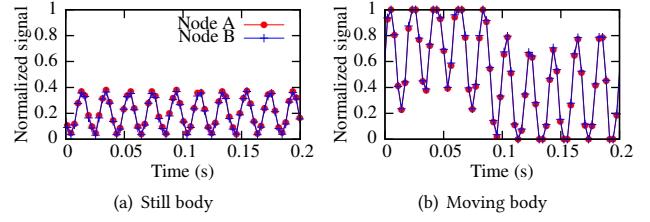


Figure 6: SEPs on a same wearer (shared ground).

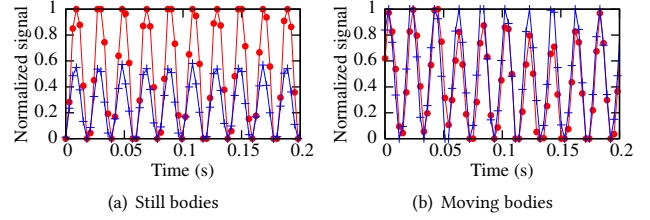


Figure 7: SEPs on different wearers (shared ground).

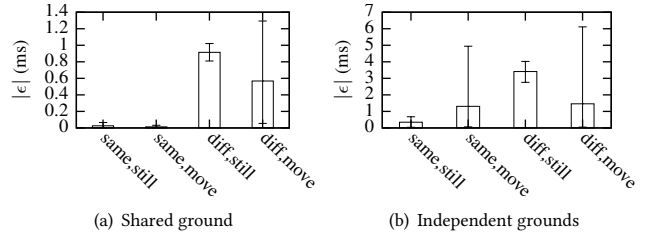


Figure 8: Absolute time displacement $|\epsilon|$ between the EMR signals captured by the two Floras in various scenarios. Error bar represents (5%, 95%) confidence interval. Each error bar is obtained from one minute of data.

Third, two researchers touch the ADC pins of the two Floras separately. Figs. 7(a) and 7(b) show the signals captured by the two nodes when the two researchers stand still and walk, respectively. The two nodes yield salient signals with different amplitudes. We note that several factors may affect the reception of powerline radiation, e.g., human body size, position and facing of the body in the electromagnetic field generated by the powerlines.

We evaluate the synchronism between the signals captured by the two Floras shown in Figs. 6 and 7. We condition the signals by first applying a band-pass filter (BPF) to remove the direct current (DC) component that may fluctuate as seen in Fig. 6(b) and then detect the zero crossings (ZCs) of the filtered signals. More details of the BPF and ZC detection will be presented in Section 5.2. We use the *time displacement* between the two signals' ZCs as the metric to evaluate their synchronism. Specifically, the time displacement, denoted by ϵ , is given by $\epsilon = t_A^{ZC} - t_B^{ZC}$, where t_A^{ZC} and t_B^{ZC} represent the ground-truth times of Node A's ZC and the corresponding ZC at Node B, respectively. Fig. 8(a) shows the error bars for $|\epsilon|$, which correspond to the scenarios in Figs. 6(a), 6(b), 7(a), and 7(b),

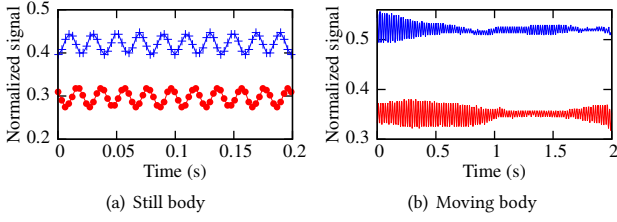


Figure 9: SEPs on a same wearer (independent grounds).

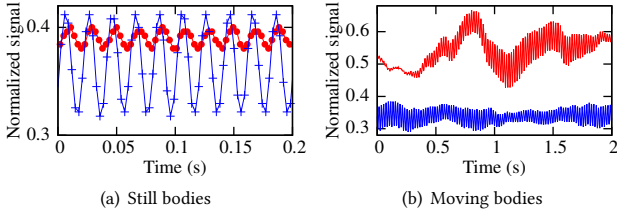


Figure 10: SEPs on different wearers (independent grounds).

respectively. In Fig. 8(a), “same” and “diff” mean the same wearer and different wearers, respectively; “still” and “move” mean standing still and walking, respectively. On the same wearer, the SEPs captured by the two Floras are highly synchronous, with an average $|\epsilon|$ of $0.9 \mu\text{s}$. On different wearers, the $|\epsilon|$ increases to about 1 ms. When the two wearers move, the average $|\epsilon|$ is 0.35 ms smaller than that when they stand still. This small difference may be caused by several affecting factors discussed earlier, i.e., human body size and etc. The human body movements increase the variance of $|\epsilon|$, since they create more signal dynamics as seen in Fig. 6(b).

4.3.2 Independent grounds. Then, we remove the connection between the two Floras’ ground pins, such that they have independent grounds. This setting is consistent with real scenarios, where the wearables are generally not wired. We conduct three tests.

Fig. 5(b) shows the two Floras’ signals when they have no physical contact with any human body. The signals have small oscillations with a frequency of 50 Hz.

Figs. 9(a) and 9(b) show the signals of the two Floras worn on two wrists of a researcher when he stands still and walks, respectively. Compared with the results in Fig. 6(a) based on a shared ground, the two signals in Fig. 9(a) have an offset in their values. This offset is the difference between the electric potentials at the two Floras’ grounds. Fig. 9(b) shows the signals over two seconds that contain about 100 SEP cycles to better illustrate the changing signal envelopes over time due to the human body movements. Compared with Fig. 6(b), the two signals in Fig. 9(b) have different signal envelopes. This is because the electric potentials at the two Floras’ grounds, which are also induced by the powerline radiation, are not fully correlated in the presence of human body movements.

Figs. 10(a) and 10(b) show the signals of the two Floras worn by two researchers each when they stand still and walk, respectively. Salient EMR signals can be observed. Moreover, the human

movements cause significant fluctuations of DC lines and the signal amplitudes, as seen in Fig. 10(b).

We also evaluate the synchronism between the two Floras’ signals. Fig. 8(b) shows the time displacement’s error bars that correspond to the scenarios in Figs. 9(a), 9(b), 10(a), and 10(b). The average $|\epsilon|$ is below 3 ms. Compared with the results in Fig. 8(a) that are based on a shared ground, the time displacements increase. This is because of the additional uncertainty introduced by the independent floating grounds of the two Floras. Nevertheless, on the same wearer, the average $|\epsilon|$ is about 1 ms only. The body movements increase the 95%-percentile of $|\epsilon|$ to 6 ms. In Section 5.2, we will use a phase-locked loop to reduce the variations of ϵ .

4.3.3 Summary. From the above measurements, we obtain the following three key observations. First, the human body can act as an antenna that effectively improves the powerline radiation reception. Second, during the human body movements, the SEP amplitude changes. However, the synchronism between the two SEP signals captured by the two nodes on the same wearer or different wearers is still acceptably preserved. In Section 5.2, we will condition the SEP signals to improve the synchronism. Third, the floating ground of a node introduces additional uncertainty, because the powerline radiation can also generate a varying electric potential at the ground pin. However, the floating ground does not substantially degrade the synchronism between the two nodes’ SEP signals. All experiments in the rest of this paper are conducted under the floating ground setting. The above three observations suggest that the SEP induced by powerline radiation is a good periodic signal that can be exploited for synchronizing wearables.

We note that, all the above measurements are conducted in a computer science laboratory that draws electricity from a single power grid phase. Thus, the SEPs received by the Floras in the laboratory have the same phase. Typically, a small area (e.g., a room and an office floor) is supplied by the same power grid phase. Thus, the wearables in the area will sense synchronized SEPs. Two remote wearables may sense different power grid phases. The voltage phase difference will become part of the synchronization error. It is 6.7 ms and 5.6 ms in 50 Hz and 60 Hz grids, respectively. We will observe this in our experiments presented in Section 7.2.

The powerline radiation and SEP are generally unavailable outdoors. Our extensive evaluation in Section 7 will show the pervasive availability of SEPs in indoor environments. As most of our lifetime is indoors (e.g., 87% on average for Americans [17]), the SEP will be available for synchronizing wearables.

5 DESIGN OF TOUCHSYNC

In this section, we present the design of TouchSync. Section 5.1 overviews the workflow of TouchSync. Section 5.2 presents the signal processing algorithms to generate stable, periodic, and synchronous impulses trains (i.e., Dirac combs) from the SEP signals. Section 5.3 presents a synchronization protocol assisted with the Dirac combs. Section 5.4 solves the integer ambiguity problem to complete synchronization.

5.1 TouchSync Workflow

TouchSync synchronizes the clock of a slave to that of a master. This paper focuses on the synchronization between a slave-master pair,

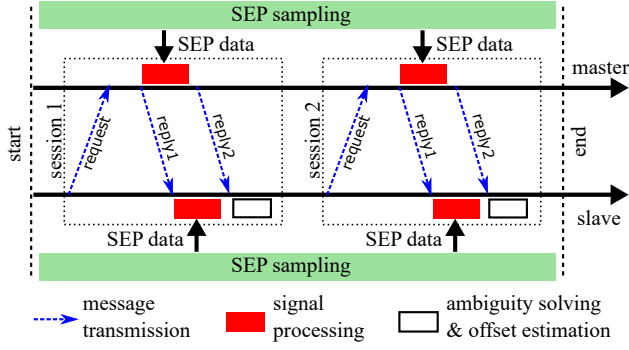


Figure 11: A synchronization process of TouchSync.

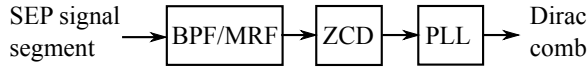


Figure 12: SEP signal processing pipeline.

which is the basis for synchronizing a network of nodes. A *synchronization process*, as illustrated in Fig. 11, is performed periodically or in an on-demand fashion. For instance, the wearer(s) may push some buttons on two wearables to start a synchronization process. The period of the synchronization can be determined by the needed clock accuracy and the clock drift rate. During the synchronization process, both the slave and the master continuously sample the SEP signals and store the timestamped samples into their local buffers. At the beginning of the synchronization process, the slave node sends a message to the master to signal the start of the sensor sampling. A synchronization process has multiple *synchronization sessions*. Fig. 11 shows two sessions. In each session, the slave and the master exchange three messages: request, reply1, and reply2. The request and reply1 are used to measure the communication delays. After transmitting the reply1, the master retrieves a segment of SEP signal from its buffer to process and transmits the processing results using the reply2 to the slave. Upon receiving the reply1, the slave retrieves a segment of SEP signal from its buffer to process. Upon receiving the reply2, the slave tries to solve the integer ambiguity problem to estimate the offset between the slave's and the master's clocks. If the ambiguity cannot be solved, another synchronization session is initiated; otherwise, the two nodes stop sampling SEPs and the slave uses the estimated offset to adjust its clock and complete the synchronization process.

5.2 SEP Signal Processing

This section presents TouchSync's signal processing illustrated as the filled blocks in Fig. 11. The objective is to generate a highly stable, periodic, and synchronous Dirac comb from a SEP signal with fluctuating DC component and jitters as shown in Figs. 10(b) and 8(b). The algorithms should be compute- and storage-efficient.

We apply a signal processing pipeline illustrated in Fig. 12. It has three steps:

Band-pass filter (BPF) or mean removal filter (MRF): We apply a 5th-order/6-tap infinite impulse response (IIR) BPF with steep boundaries of a (45 Hz, 55 Hz) passband to remove the fluctuating

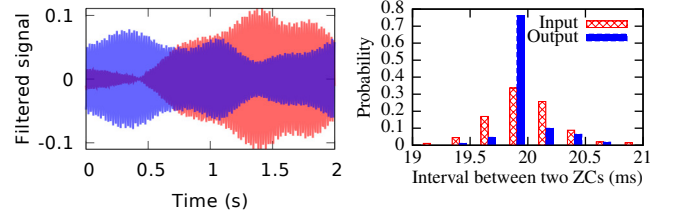


Figure 13: BPF output (each color stands for the same node in Fig. 10).

Figure 14: Jitter before and after PLL.

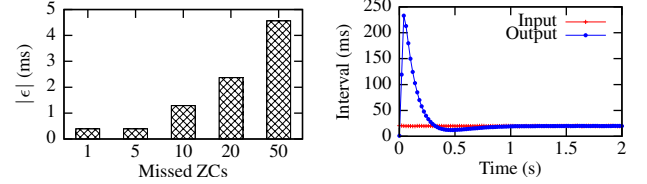


Figure 15: PLL robustness.

Figure 16: PLL convergence.

DC component and high-frequency noises of the SEP signal. The red and blue curves in Fig. 13 are the filtering results for the red and blue signals shown in Fig. 10(b). We can see that the DC components have been removed. For too resource-limited wearables, a MRF that subtracts the running average from the original signal can be used instead of the BPF for much lower compute and storage complexities. Its effect is similar to a low-pass filtering.

Zero crossing detector (ZCD): It detects the ZCs, i.e., the time instants when the filtered SEP signal changes from negative to positive. It computes a linear interpolation point between the negative and the consequent positive SEP samples as the ZC to mitigate the impact of low time resolution due to a low SEP sampling rate.

Phase-locked loop (PLL): We apply a software PLL to deal with the ZC jitters and miss detection caused by significant dynamics of the SEP signal. The PLL generates an impulse train using a loop and uses an active proportional integral (PI) controller to tune the interval between two consecutive impulses according to the time differences between the past impulses and the input ZCs. The controller skips the time differences larger than 25 ms to deal with ZC miss detection. Fig. 14 shows the distributions of the interval between two consecutive ZCs of the PLL's input and output, respectively. The PLL reduces jitters. Before PLL, the minimum and maximum intervals are 13.5 ms and 30.8 ms, respectively. After PLL, the minimum and maximum intervals are 19.4 ms and 20.7 ms, respectively. To understand PLL's robustness against ZC miss detection, we artificially drop a number of continuous ZCs to simulate an outage period and evaluate the time displacement ϵ between the PLL's output and the dropped ZCs. Fig. 15 shows the $|\epsilon|$ versus the number of missed ZCs. When 50 ZCs that last for one second are missed, the caused $|\epsilon|$ is 4.5 ms only. Moreover, after the outage period, the $|\epsilon|$ restores to sub-ms with five input ZCs only.

Though the above three steps are standard signal processing techniques, they are crucial for TouchSync. Moreover, we tune them to better cater into our needs. In [28], the PLL is also used to reduce jitters of a clock calibration signal. Different from [28] that

continuously runs PLL for continuous clock calibration, TouchSync samples SEP and runs PLL only when clock synchronization is needed. Thus, we configure the PLL to have a short convergence time of about one second, as shown in Fig. 16. Moreover, owing to a special consideration in the design of TouchSync that will be presented in Section 5.3, the signal processing algorithms operate in an “offline” manner, in that they start to work until the whole SEP signal segment to be processed becomes available. This largely simplifies the implementation of these algorithms.

5.3 NTP Assisted with Dirac Combs

TouchSync uses the synchronous Dirac combs at the slave and the master to achieve clock synchronization through multiple synchronization sessions. This section presents the protocol and the analysis for a single synchronization session.

5.3.1 Protocol for a synchronization session. A synchronization session of TouchSync is illustrated in Fig. 17. We explain it from the following two aspects.

Message exchange and timestamping: The session consists of the transmissions of three messages: request, reply1, and reply2. The request and reply1 messages are similar to the two UDP packets used by NTP. Their transmission and reception timestamps, i.e., t_1 , t_2 , t_3 , and t_4 , are obtained upon the corresponding messages are passed/received to/from the OS, as illustrated in Fig. 1(b). The master will transmit the auxiliary reply2 message to convey the results of its signal processing, which is detailed below.

Signal processing and clock offset estimation: After the master has transmitted the reply1 message, the master (i) retrieves from its signal buffer a SEP signal segment that covers the time period from t_2 to t_3 with some safeguard ranges before t_2 and after t_3 , (ii) feeds the signal processing pipeline in Section 5.2 with the retrieved SEP signal segment to produce a Dirac comb as illustrated in Fig. 17, and (iii) identifies the last impulses (LIs) in its Dirac comb that are right before the time instants t_2 and t_3 , respectively. The LIs are illustrated by thick red arrows in Fig. 17. Then, the master computes the elapsed times from t_2 's LI to t_2 and t_3 's LI to t_3 , which are denoted by ϕ_2 and ϕ_3 , respectively. The ϕ_2 and ϕ_3 are the phases of the t_2 and t_3 with respect to the Dirac comb. After that, the master transmits the reply2 message that contains t_2 , t_3 , ϕ_2 , and ϕ_3 to the slave. After receiving the reply1 message, the slave retrieves a SEP signal segment that covers the time period from t_1 to t_4 with some safeguard ranges, executes the signal processing pipeline, identifies the LIs right before t_1 and t_4 , and computes the phases ϕ_1 and ϕ_4 , as illustrated in Fig. 17. After receiving the reply2 message, based on $\{t_1, t_2, t_3, t_4\}$ and $\{\phi_1, \phi_2, \phi_3, \phi_4\}$, the slave uses the approach in Section 5.3.2 to analyze the offset between the slave's and master's clocks. From the PLL convergence speed shown in Fig. 16, we set the safeguard to one second.

Now, we discuss several design considerations for the protocol described above. TouchSync uses the request and reply1 messages to measure the clock offset, while the reply2 is an auxiliary message to convey the timestamps t_2 , t_3 and the measurements ϕ_2 , ϕ_3 . With this auxiliary message, we can decouple the task of

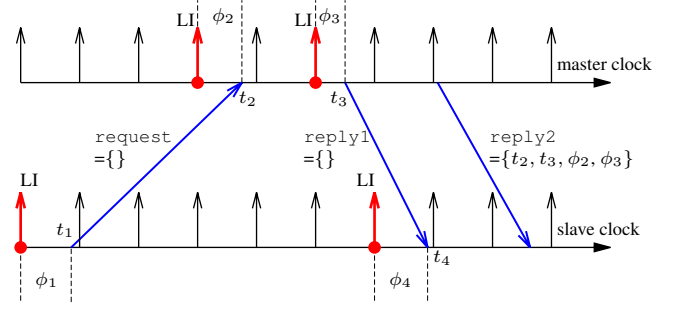


Figure 17: A synchronization session of TouchSync. The vertical arrows represent the impulses of the Dirac combs generated from the SEP signal.

timestamping the reception of request and the transmission of reply1 from the signal processing task of generating the Dirac comb and computing ϕ_2 and ϕ_3 . On many platforms (e.g., Android Wear and watchOS), continuously sampled sensor data is passed to the application block by block. With the decoupling, the master can compute ϕ_2 and ϕ_3 after the reply1 is transmitted and the needed SEP data blocks become available. This is why the signal processing algorithms in Section 5.2 can operate in an “offline” manner.

5.3.2 Clock offset analysis. We now analyze the offset between the slave's and the master's clocks based on $\{t_1, t_2, t_3, t_4\}$ and $\{\phi_1, \phi_2, \phi_3, \phi_4\}$. Denote by T the period of the Dirac comb. In our region served by a 50 Hz grid, the nominal value for T is 20 ms. To capture the small deviation from the nominal value, T can be also easily computed as the average interval between consecutive impulses of the Dirac comb. We define the *rounded phase differences* θ_q and θ_p (which correspond to the request and reply1 messages, respectively) as

$$\theta_q = \begin{cases} \phi_2 - \phi_1, & \text{if } \phi_2 - \phi_1 \geq 0; \\ \phi_2 - \phi_1 + T, & \text{otherwise.} \end{cases} \quad (1)$$

$$\theta_p = \begin{cases} \phi_4 - \phi_3, & \text{if } \phi_4 - \phi_3 \geq 0; \\ \phi_4 - \phi_3 + T, & \text{otherwise.} \end{cases} \quad (2)$$

As ϕ_k is the elapsed time from t_k 's LI to t_k , we have $0 \leq \phi_k < T$, for $k \in [1, 4]$. From Eqs. (1) and (2), we can verify that $0 \leq \theta_q < T$ and $0 \leq \theta_p < T$. From our measurements in Section 4.2, the times for transmitting the request and reply1 messages can be longer than T . Thus, we use i to denote the non-negative integer number of the Dirac comb's periods elapsed from the time of sending request to the time of receiving it at the master, and j to denote the non-negative integer number of the Dirac comb's periods elapsed from the time of sending reply1 to the time of receiving it at the slave.

We denote τ_q and τ_p the actual times for transmitting the request and the reply messages, respectively. Thus,

$$\tau_q = \theta_q + i \cdot T - \epsilon, \quad \tau_p = \theta_p + j \cdot T + \epsilon, \quad (3)$$

where ϵ is the time displacement between the slave's and master's Dirac combs. Here, we assume a constant ϵ to simplify the discussion. Therefore, the RTT computed by $\text{RTT} = (t_4 - t_1) - (t_3 - t_2)$ must satisfy

$$\text{RTT} = \tau_q + \tau_p = \theta_q + \theta_p + (i + j) \cdot T. \quad (4)$$

In Eq. (4), RTT , θ_q , and θ_p are measured in the synchronization session illustrated in Fig. 17; i and j are unknown non-negative integers. If the i and j can be determined, the estimated offset between the slave's and the master's clocks, denoted by δ , can be computed by either one of the following formulas:

$$\delta = t_1 - (t_2 - \tau_q) = t_1 - t_2 + \theta_q + i \cdot T - \epsilon, \quad (5)$$

$$\delta = t_4 - (t_3 + \tau_p) = t_4 - t_3 - \theta_p - j \cdot T - \epsilon. \quad (6)$$

It can be easily verified that the above two formulas give the same result. The analysis in the rest of this paper chooses to use Eq. (6). The ϵ is generally unknown. If we ignore it in Eq. (6) to compute δ , it becomes part of the clock offset estimation error.

Eq. (4) is an *integer-domain* underdetermined problem. Clearly, from Eq. (4), both i and j belong to the range $\left[0, \frac{\text{RTT} - \theta_q - \theta_p}{T}\right]$. Thus, Eq. (4) has a finite number of solutions for i and j . Note that, under the original NTP principle, we have a *real-domain* underdetermined problem of $\text{RTT} = \tau_q + \tau_p$ that has infinitely many solutions. NTP chooses a solution by assuming $\tau_q = \tau_p$, which does not hold in general. Thus, by introducing the Dirac combs, the *ambiguity* in determining τ_q , τ_p , and δ is substantially reduced from infinitely many possibilities to finite possibilities. Though we still have ambiguity in the integer domain, our analysis and extensive numeric results in Section 5.4 show that the ambiguity can be solved.

Note that, in [27], the periodic and synchronous power grid voltage signals collected directly from power outlets are used to synchronize two nodes that have high-speed wired network connections. The approach in [27] also uses the elapsed times from LIs (i.e., $\phi_1, \phi_2, \phi_3, \phi_4$) to deal with asymmetric communication delays and improve synchronization accuracy. However, due to the high-speed connectivity, it only considers the case where both i and j are zero. In contrast, with wireless connectivity, i and j are random and often non-zero due to the access time (cf. Section 3.1). Estimating i and j is challenging and it is the subject of Section 5.4.

5.4 Integer Ambiguity Solver (IAS)

Before we present the approach to solving integer ambiguity, we make the following two assumptions for simplicity of exposition. First, we assume that the ground-truth clock offset δ_{GT} is a constant during a synchronization process. From our performance evaluation in Section 7, TouchSync generally takes less than one second to achieve synchronization. Typical crystal oscillators found in MCUs and personal computers have drift rates of 30 to 50 ppm [10]. Thus, the maximum drift of the offset between two clocks during one second is $50 \text{ ppm} \times 1 \text{ s} \times 2 = 0.1 \text{ ms}$. This drift is smaller than the ms-level time displacement ϵ between two SEP signals, which dominates the synchronization error of TouchSync. Second, we assume $\epsilon = 0$. In Section 5.4.4, we will discuss how to deal with non-zero and time-varying ϵ .

We let i_{\min} and i_{\max} denote the minimum and maximum possible values for i ; j_{\min} and j_{\max} denote the minimum and maximum possible values for j . For instance, from our one-way message transmission time measurements summarized in Fig. 4(a), the BLE's slave-to-master transmission times are always greater than 30 ms. Thus, we may set $i_{\min} = 1$, since in our region T is 20 ms. When we have no prior knowledge about the ranges for i and j , we may simply set $i_{\min} = j_{\min} = 0$ and $i_{\max} = j_{\max} = \frac{\text{RTT} - \theta_q - \theta_p}{T}$. Section 5.4.4

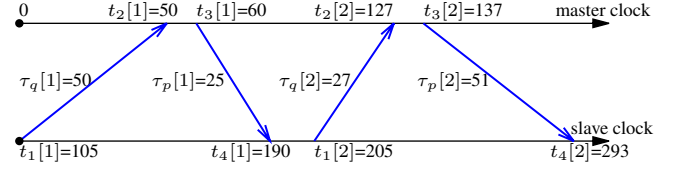


Figure 18: An example of solving the integer ambiguity. The transmissions of the auxiliary reply2 messages are omitted in the illustration.

will discuss how the use of the prior knowledge impacts on the integer ambiguity solving.

TouchSync performs multiple synchronization sessions to solve the integer ambiguity problem. In this section, we use $x[k]$ to denote a quantity x in the k^{th} synchronization session. For instance, $\text{RTT}[k]$ denotes the measured RTT in the k^{th} session. From Eqs. (4) and (6), for the k^{th} synchronization session, we have

$$\begin{cases} \text{RTT}[k] = \theta_q[k] + \theta_p[k] + (i[k] + j[k]) \cdot T; \\ \delta = t_4[k] - t_3[k] - \theta_p[k] - j[k] \cdot T; \\ i_{\min} \leq i[k] \leq i_{\max}, \quad j_{\min} \leq j[k] \leq j_{\max}. \end{cases} \quad (7)$$

If TouchSync performs K synchronization sessions, we have an underdetermined system of $2K$ equations with $(2K + 1)$ unknown variables (i.e., δ and $\{i[k], j[k] | k \in [1, K]\}$). In the integer domain, such an underdetermined system can have a unique solution.

5.4.1 An example of unique solution. We use an example in Fig. 18 to illustrate. The unit for time is ms, which is omitted in the following discussion for conciseness. In this example, $T = 20$, $i_{\min} = j_{\min} = 1$, $i_{\max} = j_{\max} = 4$, and the ground-truth clock offset $\delta_{GT} = 105$. Two synchronization sessions are performed. The timestamps and the actual message transmission delays are shown in Fig. 18. The ground-truth values for i and j in the two synchronization sessions are: $i[1] = 2$, $j[1] = 1$, $i[2] = 1$, and $j[2] = 2$. The RTTs can be computed as $\text{RTT}[1] = 75$ and $\text{RTT}[2] = 78$. With any synchronous Dirac combs, from Eqs. (1) and (2), the rounded phase differences computed by the two nodes must be $\theta_q[1] = 10$, $\theta_p[1] = 5$, $\theta_q[2] = 7$, and $\theta_p[2] = 11$. For the first synchronization session, Eq. (7) has two possible solutions only:

$$\{i[1]=1, j[1]=2, \delta=85\}, \quad \{i[1]=2, j[1]=1, \delta=105\}. \quad (8)$$

For the second synchronization session, Eq. (7) has two possible solutions only as well:

$$\{i[2]=1, j[2]=2, \delta=105\}, \quad \{i[2]=2, j[2]=1, \delta=125\}. \quad (9)$$

From Eqs. (8) and (9), $\delta = 105$ is the only common solution. Thus, we conclude that δ must be 105.

5.4.2 Program for solving integer ambiguity. From the above example, due to the diversity of the ground-truth values of i and j in multiple synchronization sessions, the intersection of the δ solution spaces of these synchronization sessions can be a single value. Thus, the integer ambiguity problem is solved. On the contrary, if the ground-truth i and j do not change over multiple synchronization sessions, the ambiguity remains. As explained in Section 3.1, with application-layer timestamping, the message transmission times are highly dynamic due to the uncertain OS overhead and MAC.

Such uncertainties and dynamics, which are undesirable in the original theme of NTP, interestingly, become desirable for solving the integer ambiguity in TouchSync.

From the above key observation, TouchSync performs the synchronization session illustrated in Fig. 17 *repeatedly* until the intersection among the δ solution spaces of all the synchronization sessions converges to a single value. Algorithms 1 and 2 provide the pseudocode for TouchSync's slave and master programs.

Algorithm 1 Slave's pseudocode for a synchronization process

```

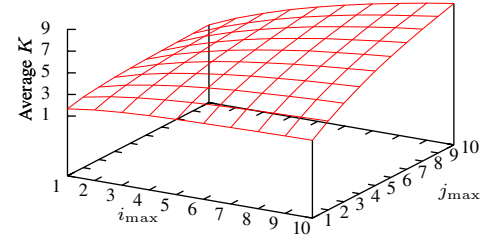
1: Global variables:  $t_1, t_4, \delta$ 's solution space  $\Delta = \emptyset$ , session index  $k = 0$ 
2:
3: command start_sync_session() do
4:    $k = k + 1$ 
5:    $t_1 = \text{read\_system\_clock}()$ 
6:   send message request = {} to master
7: end command
8:
9: event reply1 received from master do
10:   $t_4 = \text{read\_system\_clock}()$ 
11:  wait until SEP data covering  $t_1$  and  $t_4$  are available
12:  run the SEP signal processing pipeline in Section 5.2
13:  compute  $\phi_1$  and  $\phi_4$ 
14: end event
15:
16: event reply2 received from master do
17:  compute  $\theta_q$  and  $\theta_p$  using Eqs. (1) and (2), respectively
18:   $\text{RTT} = (t_4 - t_1) - (\text{reply2}.t_3 - \text{reply2}.t_2)$ 
19:  solve Eq. (7),  $\Delta'$  denotes the set of all possible solutions for  $\delta$ 
20:  if  $k == 1$ :  $\Delta = \Delta'$ ; else:  $\Delta = \Delta \cap \Delta'$ 
21:  if  $\Delta$  has only one element  $\delta$ : use  $\delta$  to adjust clock;
22:  else: start_sync_session() // start a new synchronization session
23: end event
```

Algorithm 2 Master's pseudocode for a synchronization process

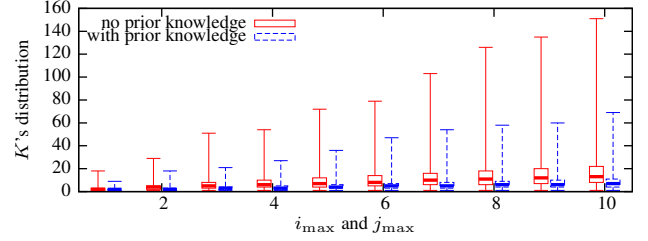
```

1: event request received from slave do
2:   $t_2 = \text{read\_system\_clock}()$ 
3:  ... // execute other compute tasks
4:   $t_3 = \text{read\_system\_clock}()$ 
5:  send message reply1 = {} to slave
6:  wait until SEP data covering  $t_2$  and  $t_3$  are available
7:  run the SEP signal processing pipeline in Section 5.2
8:  compute  $\phi_2$  and  $\phi_3$ 
9:  send message reply2 =  $\{t_2, t_3, \phi_2, \phi_3\}$  to slave
10: end event
```

5.4.3 Convergence speed. We run a set of numeric experiments to understand the convergence speed of the IAS. We use the number of synchronization sessions until convergence to characterize the convergence speed, which is denoted by K in the rest of this paper. We fix i_{\min} and j_{\min} to be zero. For a certain setting $\langle i_{\max}, j_{\max} \rangle$, we conduct 100,000 synchronization processes to assess the distribution of K . For each synchronization session of a synchronization process, we randomly and uniformly generate the ground-truth i and j , as well as θ_q and θ_p within their respective ranges, i.e., $i \in [0, i_{\max}]$, $j \in [0, j_{\max}]$, and $\theta_q, \theta_p \in [0, T)$. Then, we simulate the integer ambiguity solving program presented in Section 5.4.2 to measure



(a) Average K vs. i_{\max} and j_{\max} .



(b) Box plot for K (box represents the 1st and 3rd quartiles; whiskers represent minimum and maximum values.)

Figure 19: Convergence speed of IAS.

the K for each synchronization process. In practice, the i, j, θ_q and θ_p may not follow the uniform distributions. But the numeric results here help us understand the convergence speed. In Section 7.2, we will evaluate the convergence speed in real-world settings.

In Fig. 19(a), each grid point is the average of all K values in the 100,000 synchronization processes under a certain $\langle i_{\max}, j_{\max} \rangle$ setting. Fig. 19(b) shows the box plot for K under each setting where $i_{\max} = j_{\max}$. We note that all simulated synchronization processes converge. From the two figures, even if $i_{\max} = j_{\max} = 10$ (which means that the one-way communication delays are up to 200 ms for $T = 20$ ms), the average K is nine only. Although the K 's distribution has a long tail as shown in Fig. 19(b), 75% of the K values are below 11. This result is consistent with our real experiment results in Table 2 of Section 7.2, where most K values are two only and the largest K is 12.

5.4.4 Discussions. First, we discuss how to address non-zero and time-varying ϵ . From the analysis in Eq. (7) that is based on $\epsilon = 0$, the difference between two δ solutions is multiple of T . This can also be seen from Eqs. (8) and (9). In practice, ϵ can be non-zero and time-varying. It will be a major part of the δ estimation error. From Fig. 8(b), the $|\epsilon|$ is at most 6 ms. Thus, the resulted variation to the δ solutions will be less than a half and one third of T , in the regions served by 60 Hz and 50 Hz power grids, respectively. Therefore, we can still correctly identify the correspondence among the δ elements in the set intersection operation in Line 20 of the slave's program in Algorithm 1. Specifically, if two δ elements have a difference smaller than $T/2$, they should be considered the same element in the set intersection operation; otherwise, they are different elements. For this correspondence identification to be correct, the ϵ needs to be smaller than $T/2$. After convergence, the final δ can be computed as the average of the δ elements that are considered the same. We have incorporated this in our implementation of TouchSync.

Second, we discuss how the use of the prior knowledge (i.e., i_{\min} , i_{\max} , j_{\min} , and j_{\max}) impacts on the integer ambiguity solving. With the prior knowledge, we may shrink the search range for i and j to speed up the convergence of the IAS. The prior knowledge can be based on the statistical information obtained in offline experiments. For instance, a group of the box plots are the results for the IAS with the prior knowledge of i_{\max} and j_{\max} . The IAS can search the i and j within the ranges of $[0, i_{\max}]$ and $[0, j_{\max}]$, respectively. The other group of results are for the IAS without the prior knowledge. Thus, the IAS has to search within the range of $\left[0, \frac{RTT - \theta_q - \theta_p}{T}\right]$ for both i and j . We can see that, if no prior knowledge is used, the K increases. But the IAS still always converges. Once the IAS converges, the synchronization error of TouchSync mainly depends on the time displacement ϵ .

6 IMPLEMENTATION OF TOUCHSYNC

6.1 touchsync.h Header and Implementations

Designed as an application-layer clock synchronization approach, TouchSync can be implemented as an app or part of an app, purely based on the standard wearable OS calls to sample the SEP signal, exchange network messages, and timestamp them in the application layer. To simplify the adoption of TouchSync by application developers, we have implemented TouchSync's platform-independent tasks (i.e., SEP signal processing and IAS) in ANSI C and provide them in a `touchsync.h` header file [31]. As most embedded and IoT platforms are C compatible, our C implementation is applicable to a wide range of wearables. Wrappers for other programming languages can also be implemented. The header file defines a circular buffer to store the SEP signal. It provides four functions to be used by the application developer:

- `buf_add()` pushes a new SEP sample to the circular buffer;
- `on_receive_reply1()` implements Line 11-13 of Algorithm 1;
- `on_receive_reply2()` implements Line 17-20 of Algorithm 1;
- `send_reply1_done()` implements Line 6-8 of Algorithm 2.

Other tasks of TouchSync, i.e., sensor sampling, synchronization message exchange and timestamping, are platform dependent. We leave them for the application developer to implement. As these tasks are basics for embedded programming, by using the four functions provided by `touchsync.h`, application developers without much knowledge in signal processing can readily implement TouchSync on different platforms. Our own Arduino and TinyOS programs that implement TouchSync's work flow have about 50 and 150 lines of code only, respectively.

6.2 Benchmarking

To understand the overhead of TouchSync, we deploy our TinyOS and Arduino implementations to Zolertia's Z1 motes [13] and Floras, respectively. The Z1 mote is equipped with an MSP430 MCU (1 MHz, 8 KB RAM) and a CC2420 802.15.4 radio. Both implementations sample SEP at 333 Hz. On Z1, we configure the length of the circular buffer defined in `touchsync.h` to be 512. Thus, this circular buffer can store 1.5 seconds of SEP data. This is sufficient, because the time periods $[t_2, t_3]$ and $[t_1, t_4]$ that should be covered by the SEP signal segments to be retrieved from the circular buffer and

Table 1: Storage and compute overhead of TouchSync.

Platform	Memory use (KB)		Processing time (ms)			
	ROM	RAM	BPF/MRF	ZCD	PLL	IAS
Z1	10	5	364	9	48	1
Flora	17	1.9*	1.3	3	15	0.8

* Estimated based on buffer lengths.

processed by the master and slave are generally a few ms and below 100 ms, respectively. On Flora, we configure the circular buffer length to be 400 and redefine its data type such that TouchSync can fit into Flora's limited RAM space of 2.5 KB. Table 1 tabulates the memory usage of TouchSync and the computation time of different processing tasks. On Z1, a total of 421 ms processing time is needed for a synchronization session. The BPF uses a major portion of the processing time. Flora cannot adopt BPF because of RAM shortage. It uses MRF instead, which consumes much less RAM and processing time.

Z1 and Flora are two representative resource-constrained platforms. The successful implementations of TouchSync on them suggest that TouchSync can also be readily implemented on other more resourceful platforms.

7 PERFORMANCE EVALUATION

We conduct extensive experiments to evaluate the performance of TouchSync in various real environments. Each experiment uses two Flora nodes, which act as the TouchSync slave and master, respectively. As Flora does not support BLE master mode, the two Floras cannot communicate directly. Thus, we use a RPi that operates as a BLE master to relay the data packets between the two Floras. This setting is consistent with most body-area networks with a smartphone as the hub. If the hub can also sample powerline radiation or SEP¹, each wearable can also synchronize with the hub directly using TouchSync. We use the approach discussed in Section 4.1 to obtain the ground truth clock of each Flora. The details and the results of our experiments are presented below.

7.1 Signal Strength and Wearing Position

As the intensity of powerline radiation attenuates with distance, SEPs will have varying signal strength. Thus, we evaluate the impact of the SEP signal strength on the performance of the signal processing pipeline in Section 5.1. We measure the signal strength as follows. For a full-scale sinusoid signal with a peak-to-peak amplitude of one (normalized using ADC's reference voltage), its standard deviation is $0.5/\sqrt{2} = 0.354$. The signal strength of a normalized sinusoid with a standard deviation of σ is defined as $\sigma/0.354$. Thus, a 100% signal strength suggests a full-scale signal for the ADC. For this experiment, we use a Flora to record a SEP signal. The strength of this signal is 34%. We feed the signal processing pipeline with this signal to generate a series of *baseline* ZCs. Then, we scale down the amplitude of this signal, re-quantize it, and process it using the pipeline to generate another series of ZCs. We use the mean absolute error (MAE) of these ZCs with respect to the

¹Our preliminary experiments show that a smartphone can capture the powerline radiation by sampling its built-in microphone and then applying a BPF on the collected audio data. Thus, it is possible to implement TouchSync on smartphones.

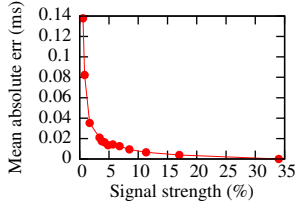


Figure 20: Impact of signal strength.

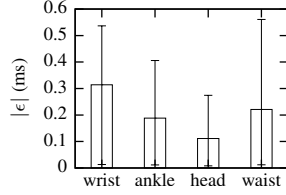


Figure 21: Impact of wearing position.

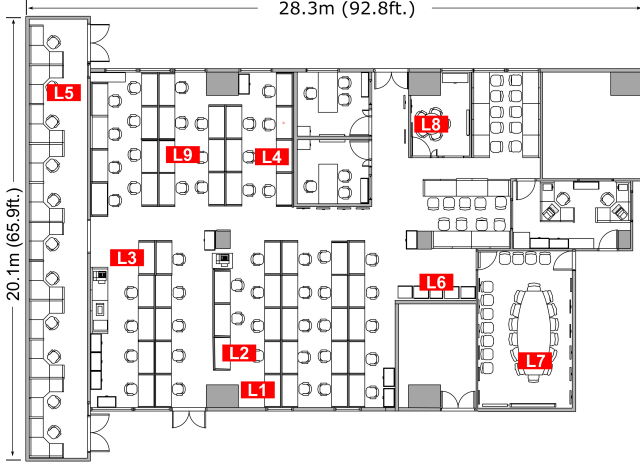


Figure 22: Laboratory floor plan with test points marked.

baseline ZCs as the error metric. Fig. 20 shows the MAE versus the strength of the scaled down signal. When we scale down the signal by 60 times, yielding a signal strength of 0.6%, the ZCs' MAE is 0.14 ms only. This suggests that TouchSync can still detect the ZCs accurately even when the SEP signal is rather weak. TouchSync can inform the wearer if the signal is too weak to detect ZCs reliably.

We evaluate the impact of the wearing position on the synchronism of SEP signals. A researcher wears a Flora on his left wrist. Then, he conducts four tests by fixing the second Flora to his right wrist, right ankle, forehead, and waist, respectively. Each test lasts for two minutes. Fig. 21 shows the error bars (5%-95% confidence interval) for the absolute time displacement $|\epsilon|$ between the two Floras in these four tests. The average $|\epsilon|$ values in the four tests are close. This suggests that the wearing positions have little impact on the synchronism of SEP signals and the synchronization accuracy of TouchSync.

7.2 Evaluation in Various Environments

We evaluate the SEP signal strength and the accuracy of TouchSync in various indoor environments.

7.2.1 Laboratory. We conduct experiments in a computer science laboratory with about 100 seats and various office facilities (lights, computers, printers, projectors, meeting rooms, etc). Fig. 22 shows the laboratory's floor plan. We arbitrarily select nine test points, marked by "Lx" in Fig. 22. A researcher carries the Floras to

Table 2: Signal strength and TouchSync accuracy.

	Test point	Without skin contact			With skin contact		
		Signal strength	K	error (ms)	Signal strength	K	error (ms)
Laboratory	L1	2.6%	3	-0.2	84.7%	2	-0.7
	L2	3.2%	2	0.3	31.5%	3	-0.7
	L3	2.3%	2	-2.5	26.1%	2	0.5
	L4	4.0%	1	-0.6	33.7%	2	0.0
	L5	0.8%	15	1.1	3.3%	2	-0.2
	L6	5.7%	10	-0.4	39.5%	10	-0.0
	L7	2.3%	n.a.	n.a.	3.0%	2	-0.9
	L8	4.6%	2	-1.5	8.3%	2	0.6
	L9	2.6%	1	-1.2	67.4%	2	-0.9
Home	H1	4.2%	2	-1.1	8.9%	2	-0.8
	H2	3.4%	1	-0.9	14.5%	2	-1.0
	H3	4.6%	1	-1.3	44.9%	2	0.2
	H4	7.8%	n.a.	n.a.	39.2%	2	0.3
	H5	3.8%	1	-1.6	3.9%	1	-2.8
	H6	3.9%	4	-4.4	9.9%	2	-2.3
	H7	5.0%	2	-1.9	6.8%	1	-2.9
	H8	8.2%	1	-11.5	54.7%	4	-1.3
	H9	2.9%	1	-2.4	9.1%	1	-1.3
Office	O1	4.0%	n.a.	n.a.	3.3%	4	0.4
	O2	5.6%	1	-7.9	2.9%	2	-1.6
	O3	1.7%	1	-0.4	3.9%	2	-0.3
	O4	5.4%	3	-2.5	5.8%	2	-0.8
	O5	4.8%	6	-6.2	5.6%	12	-0.2
Corridor	C1	3.6%	12	0.1	4.4%	11	0.7
	C2	6.2%	2	0.6	44.2%	2	-1.0
	C3	5.8%	1	-7.6	4.4%	1	-1.1
	C4	1.9%	1	-6.0	2.2%	1	-2.8
	C5	1.9%	2	-3.7	5.8%	1	-1.0

* n.a. means that TouchSync cannot converge because of large ϵ .

each test point and conducts two experiments. In the first experiment, the Floras have no physical contact with human body; in the second experiment, the researcher wears the two Floras. Thus, the experiment evaluates the same-wearer scenario. The example applications mentioned in Section 1, i.e., wireless earbuds, motion analysis, and muscle activation monitoring, belong to this scenario. Each synchronization session takes about 150 ms. A synchronization process completes once the IAS converges.

The first part of Table 2 shows the SEP's signal strength, the number of synchronization sessions until convergence (K), and the clock offset estimation error at each test point. Without skin contact, the signal strength is a few percent only. But TouchSync can still achieve a 3 ms accuracy. At L7, TouchSync cannot converge because of large and varying ϵ . The skin contact significantly increases the signal strength. Moreover, TouchSync converges after two synchronization sessions in most cases. For $K = 2$, a synchronization process takes less than one second. The absolute clock offset estimation errors are below 1 ms, lower than those without skin contact. However, without skin contact, the accuracy does not substantially degrade. This suggests that, TouchSync is resilient to the loss of skin contact due to say loose wearing.

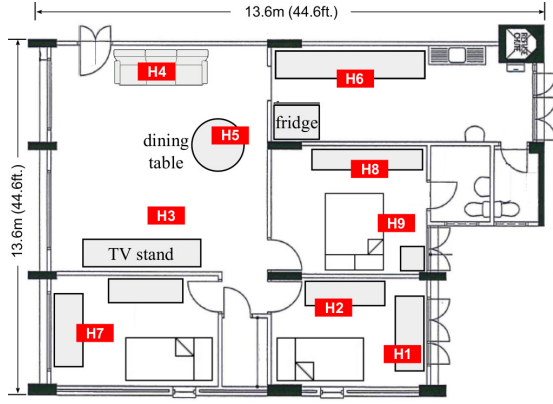


Figure 23: Home floor plan with test points marked.

7.2.2 Home. We conduct experiments in a 104 m² three-bedroom home with typical home furniture and appliances. Fig. 23 shows the home's floor plan. We arbitrarily select nine test points, marked by "Hx" in Fig. 23. The second part of Table 2 shows the results. Without skin contact, the signal strength results are similar to those obtained in the laboratory. With skin contact, the signal strength increases and the absolute clock offset estimation errors are below 3 ms.

7.2.3 Office. The third part of Table 2 shows the results obtained at five test points in a 15 m² office. At test points O1 and O2, the signal strength with skin contact is slightly lower than that without skin contact. This is possible as the two tests were conducted during different time periods and the powerline radiation may vary over time due to changed electric currents. With skin contact, TouchSync gives sub-ms accuracy except at O2.

7.2.4 Corridor. We select five test points with equal spacing in a 200 m corridor of a campus building. The fourth part of Table 2 gives the results. With skin contact, TouchSync yields absolute clock offset estimation errors of about 1 ms except at C4. The errors with skin contact are lower than those without skin contact.

7.2.5 Discussions. In summary, with skin contact, TouchSync gives sub-ms clock offset estimation errors at 20 test points out of totally 28 test points in Table 2. All errors are below 3 ms.

From Table 2, at 3 out of 28 test points, the test without skin contact gives higher signal strength than that with skin contact. This is because the two tests are conducted sequentially and the EMR may change over time. Moreover, the signal strength exhibits significant variation at different locations. This is because the EMR decays with the distance from the powerline. Nevertheless, the above results show the pervasive availability of SEP in indoor environments. In our experiments presented above, we do not observe interference from electrical appliances that leads to wrong clock synchronization. In our future work, we will conduct extensive investigation on the impact of the possible electromagnetic radiation from various electrical appliances on TouchSync.

We now discuss the energy consumption of TouchSync. Compared with NTP, TouchSync's SEP sampling and extra synchronization sessions incur additional energy consumption. SEP sampling

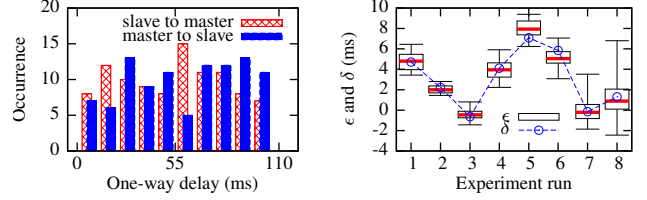


Figure 24: One-way delays over a ngrok tunnel. Figure 25: Accuracy of TouchSync-over-Internet.

is performed during the clock synchronization process only. The ADC's power consumption is much lower than the radio's. For instance, the ADC embedded in the TI MSP430G2x53 MCU consumes 0.6 mA only, much lower than the power consumption of ZigBee and Bluetooth radios which is typically tens of mA. Thus, the radio's energy consumption during the IAS's convergence process dominates TouchSync's energy consumption. From Table 2, with skin contact, the IAS converges within 2.9 synchronization sessions on average. Thus, TouchSync's energy consumption is estimated as about three times of NTP's.

7.3 TouchSync-over-Internet

Tightly synchronizing wearables over long physical distances is often desirable. For instance, in distributed virtual reality applications, tight clock synchronization among participating sensing and rendering devices that may be geographically distributed is essential [6, 9]. Although the synchronization can be performed in a hop-by-hop manner (e.g., wearables ↔ smartphone ↔ cloud), errors accumulate over hops. In particular, tightly synchronizing a smartphone to global time has been a real and challenging problem – tests showed that the synchronization through LTE and Wi-Fi experiences hundreds of ms jitters [18]. In contrast, TouchSync can perform end-to-end synchronizations for wearables distributed in a geographic region served by the same power grid. The basis is that, the 50/60 Hz power grid voltage, which generates the powerline radiation, is highly synchronous across the whole power grid [30]. TouchSync can also synchronize wearables directly with a cloud server in the same region. The smartphone merely relays the messages exchanged among the wearables and the cloud server if the wearables cannot directly access Internet. The cloud server can use a sensor directly plugged in to a power outlet to capture the power grid voltage. Owing to the Internet connectivity, the end-to-end synchronization scheme greatly simplifies the system design and implementation.

We conduct a proof-of-concept experiment of TouchSync-over-Internet as follows. Two researchers carry a Flora-RPi setup each to two buildings that are about 10 km apart. The RPi is attached with an Adafruit GPS receiver to obtain ground-truth coordinated universal time (UTC) with μ s accuracy. The two nodes, one as TouchSync master and the other as TouchSync slave, communicate through a tunnel established by ngrok 1.7, an open-source reverse proxy often adopted for creating IoT networks. Fig. 24 shows the distributions of the two one-way delays over the ngrok tunnel. We can see that the ngrok exhibits greater dynamics than the BLE link (cf. Fig. 4(a)). We evaluate TouchSync-over-Internet for eight times

during a day. Fig. 25 shows the box plots of the time displacements (i.e., ϵ) between the SEPs captured by the two nodes. We can see that ϵ varies from -2 ms to 9 ms during the day. From the building managements, the two rooms where the master and slave nodes are located draw electricity from the R and Y phases of the power grid, respectively. There is a phase difference of $20/3 = 6.7$ ms between these two phases. Moreover, from power engineering, the difference between the power grid voltage phases at different geographic locations is non-zero and time-varying. These factors lead to the non-zero and time-varying ϵ in Fig. 25. The dotted line in Fig. 25 shows the synchronization errors of TouchSync-over-Internet (i.e., δ) in various experiment runs. They are within the range of ϵ , since ϵ is the major source of TouchSync's synchronization error. The largest δ is 7 ms. The integer ambiguity solver converges within 4 to 13 synchronization sessions.

8 CONCLUSION

TouchSync synchronizes the clocks of wearables by exploiting the wearers' skin electric potentials induced by powerline radiation. Different from existing WSN clock synchronization approaches that find difficulties in being applied on diverse IoT platforms due to their need of hardware-level packet timestamping or non-trivial extra hardware, TouchSync can be readily implemented as an app based on standard wearable OS calls. Extensive evaluation shows TouchSync's synchronization errors of below 3 ms and 7 ms on the same wearer and between two wearers 10 km apart, respectively.

ACKNOWLEDGMENTS

The authors wish to thank our shepherd Dr. Lu Su and the anonymous reviewers for providing valuable feedback on this work. This research was funded by a Start-up Grant at Nanyang Technological University.

REFERENCES

- [1] Marie Chan, Daniel EstèVe, Jean-Yves Fourniols, Christophe Escriba, and Eric Campo. 2012. Smart Wearable Systems: Current Status and Future Challenges. *Artificial Intelligence in Medicine* 56, 3 (Nov. 2012), 137–156.
- [2] Yin Chen, Qiang Wang, Marcus Chang, and Andreas Terzis. 2011. Ultra-low power time synchronization using passive radio receivers. In *Proceedings of the 10th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, Chicago, IL, USA, 235–245.
- [3] Mihail C. Dinescu, Joseph Mazza, Adam Kujanski, Brian Gaza, and Michael Sagan. 2015. Synchronizing wireless earphones. (April 7 2015). <https://www.google.com/patents/US9002044> US Patent 9,002,044.
- [4] Jeremy Elson, Lewis Girod, and Deborah Estrin. 2002. Fine-grained Network Time Synchronization Using Reference Broadcasts. *SIGOPS Operating Systems Review* 36, SI (Dec. 2002), 147–163.
- [5] Raspberry Pi Foundation. 2017. Raspberry Pi 3 Model B. (2017). <https://www.raspberrypi.org/products/raspberry-pi-3-model-b/>.
- [6] Martin Friedmann, Thad Starner, and Alex Pentland. 1992. Synchronization in Virtual Realities. *Presence: Teleoperators and Virtual Environments* 1, 1 (Jan. 1992), 139–144.
- [7] Saurabh Ganeriwal, Ram Kumar, and Mani B. Srivastava. 2003. Timing-sync Protocol for Sensor Networks. In *Proceedings of the 1st International Conference on Embedded Networked Sensor Systems (SenSys)*. ACM, Los Angeles, California, USA, 138–149.
- [8] Adam Geboff, Sriram Hariharan, Joakim Linde, Li-Quan Tan, and Shahrooz Shahparnia. 2015. Device synchronization over bluetooth. (April 2 2015). <https://www.google.com/patents/US20150092642> US Patent App. 14/496,314.
- [9] Felix G. Hamza-Lup and Jannick P. Rolland. 2004. Scene Synchronization for Real-time Interaction in Distributed Mixed Reality and Virtual Reality Environments. *Presence: Teleoperators and Virtual Environments* 13, 3 (July 2004), 315–327.
- [10] Tian Hao, Ruogu Zhou, Guoliang Xing, and Matt Mutka. 2011. WizSync: Exploiting Wi-Fi Infrastructure for Clock Synchronization in Wireless Sensor Networks. In *Proceedings of the 32nd Real-Time Systems Symposium (RTSS)*. IEEE, Vienna, Austria, 149–158.
- [11] IDC. 2016. IDC Forecasts Wearables Shipments to Reach 213.6 Million Units Worldwide in 2020. (2016). <http://www.idc.com/getdoc.jsp?containerId=prUS41530816>.
- [12] IEEE. 2008. IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems. *IEEE Std 1588-2008 (Revision of IEEE Std 1588-2002)* 1, 1 (July 2008), 1–300. <https://doi.org/10.1109/IEEESTD.2008.4579760>
- [13] Zolertia Inc. 2017. Z1 mote. (2017). <http://zolertia.io/z1>.
- [14] Adafruit Industries. 2017. Adafruit FLORA. (2017). <https://www.adafruit.com/category/92>.
- [15] Adafruit Industries. 2017. Adafruit nRF51 BLE Library. (2017). <https://learn.adafruit.com/adafruit-feather-32u4-bluefruit-le/installing-ble-library>.
- [16] Adafruit Industries. 2017. Adafruit Python BluefruitLE. (2017). https://github.com/adafruit/Adafruit_Python_BluefruitLE.
- [17] Neil E. Klepeis, William C. Nelson, Wayne R. Ott, John P. Robinson, Andy M. Tsang, Paul Switzer, Joseph V. Behar, Stephen C. Hern, and William H. Engelmann. 2001. The National Human Activity Pattern Survey (NHAPS): a resource for assessing exposure to environmental pollutants. *Journal of Exposure Analysis and Environmental Epidemiology* 11, 3 (2001), 231–252.
- [18] Patrick Lazik, Niranjini Rajagopal, Bruno Sinopoli, and Anthony Rowe. 2015. Ultrasonic time synchronization and ranging on smartphones. In *Proceedings of 21st IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*. IEEE, Seattle, WA, USA, 108–118.
- [19] Liqu Li, Guoliang Xing, Limin Sun, Wei Huangfu, Ruogu Zhou, and Hongsong Zhu. 2011. Exploiting FM Radio Data System for Adaptive Clock Calibration in Sensor Networks. In *Proceedings of the 9th International Conference on Mobile Systems, Applications, and Services (MobiSys)*. ACM, Bethesda, Maryland, USA, 169–182.
- [20] Yang Li, Rui Tan, and David K. Y. Yau. 2017. Natural Timestamping Using Powerline Electromagnetic Radiation. In *Proceedings of the 16th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. ACM, Pittsburgh, Pennsylvania, 55–66.
- [21] Zhenjiang Li, Wenwei Chen, Cheng Li, Mo Li, Xiang-Yang Li, and Yunhao Liu. 2012. FLIGHT: Clock Calibration Using Fluorescent Lighting. In *Proceedings of the 18th Annual International Conference on Mobile Computing and Networking (MobiCom)*. ACM, Istanbul, Turkey, 329–340.
- [22] Konrad Lorincz, Bor-rong Chen, Geoffrey Werner Challen, Atanu Roy Chowdhury, Shyamal Patel, Paolo Bonato, and Matt Welsh. 2009. Mercury: A Wearable Sensor Network Platform for High-fidelity Motion Analysis. In *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems (SenSys)*. ACM, Berkeley, California, 183–196.
- [23] Miklós Maróti, Branislav Kusy, Gyula Simon, and Ákos Lédeczi. 2004. The Flooding Time Synchronization Protocol. In *Proceedings of the 2nd International Conference on Embedded Networked Sensor Systems (SenSys)*. ACM, Baltimore, MD, USA, 39–49.
- [24] Frank Mokaya, Roland Lucas, Hae Young Noh, and Pei Zhang. 2015. MyoVibe: Vibration Based Wearable Muscle Activation Detection in High Mobility Exercises. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*. ACM, Osaka, Japan, 27–38.
- [25] Frank Mokaya, Roland Lucas, Hae Young Noh, and Pei Zhang. 2016. Burnout: A Wearable System for Unobtrusive Skeletal Muscle Fatigue Estimation. In *Proceedings of the 15th International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, Vienna, Austria, 1–12.
- [26] NTP Project. 2017. NTP: The Network Time Protocol. (2017). <http://www.ntp.org/>.
- [27] Dima Rabadi, Rui Tan, David K.Y. Yau, and Sreejaya Viswanathan. 2017. Taming Asymmetric Network Delays for Clock Synchronization Using Power Grid Voltage. In *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security (ASIACCS)*. ACM, Abu Dhabi, United Arab Emirates, 874–886.
- [28] Anthony Rowe, Vikram Gupta, and Ragunathan (Raj) Rajkumar. 2009. Low-power Clock Synchronization Using Electromagnetic Energy Radiating from AC Power Lines. In *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems (SenSys)*. ACM, Berkeley, California, 211–224.
- [29] Symmetricom, Inc. 2011. Symmetricom announces general availability of industry's first commercially-available chip scale atomic clock. (2011). <http://www.businesswire.com/news/home/20110118005182/en/Symmetricom-Announces-General-Availability-Industry's-Commercially-Available-Chip>.
- [30] Sreejaya Viswanathan, Rui Tan, and David K.Y. Yau. 2016. Exploiting Power Grid for Accurate and Secure Clock Synchronization in Industrial IoT. In *Proceedings of the 37th IEEE Real-Time Systems Symposium (RTSS)*. IEEE, Porto, Portugal, 146–156.
- [31] Zhenyu Yan, Yang Li, Rui Tan, and Huang Jun. 2017. TouchSync implementation. (2017). <https://github.com/yanmarvin/touchsync>.