
Rapid Clock Synchronisation for Ubiquitous Sensing Services Involving Multiple Smartphones

Chu Luo

University of Melbourne
Parkville, VIC 3052, Australia
CHUL3@student.unimelb.edu.au

Henri Koski

Mikko Korhonen
University of Oulu
Oulu, Finland
{hekoski, kmikko}@ee.oulu.fi

Jorge Goncalves

University of Melbourne
Parkville, VIC 3052, Australia
Jorge.Goncalves@unimelb.edu.au

Theodoros Anagnostopoulos

Ordnance Survey, UK
{firstname.lastname}@os.uk

Shin'Ichi Konomi

University of Tokyo, Tokyo, Japan
konomi@csis.u-tokyo.ac.jp

Simon Klakegg

University of Oulu, Oulu, Finland
Simon.Klakegg@oulu.fi

Vassilis Kostakos

University of Melbourne
Parkville, VIC 3052, Australia
Vassilis.Kostakos@unimelb.edu.au

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the Owner/Author.

UbiComp/ISWC'17 Adjunct, September 11–15, 2017, Maui, HI, USA
© 2017 Copyright is held by the owner/author(s).
ACM ISBN 978-1-4503-5190-4/17/09.
<https://doi.org/10.1145/3123024.3124432>

Abstract

This paper investigates the precision of rapid clock synchronisation for ubiquitous sensing services which consist of multiple smartphones. Specifically, we consider scenarios where multiple smartphones are used to sense physical phenomena, and subsequently the sensor data from multiple distributed devices is aggregated. We observe that the accumulated clock drift for smartphones can be more than 150ms per day in the worst case. We show that solutions using the public Network Time Protocol (NTP) can be noisy with errors up to 1800ms in one request. We describe a rapid clock synchronisation technique that reduces drift to 10ms on average (measured by linear regression) and achieves pair-wise synchronisation between smartphones with an average of 27ms (measured by accelerometer), following a Gaussian-like distribution. Our results provide a lower bound for rapid clock synchronisation as a guide when developing ubiquitous sensing services using multiple smartphones.

Author Keywords

Mobile devices; smartphones; clock synchronisation; mobile sensing.

ACM Classification Keywords

H.5.m. Information interfaces and presentation.

Introduction

The rich sensors of today's smartphones (e.g. accelerometer, magnetometer, GPS) are increasingly enabling researchers to adopt smartphones for scientific and engineering sensing applications. Particularly, several approaches focus on time-critical monitoring using multiple smartphones, e.g. earthquake monitoring [1] and pothole detection [8].

However, clock synchronisation has not been adequately addressed for smartphone sensing applications. Previous work highlights that the acquisition of accurate timing is a fundamental and non-trivial task for time-critical mobile phone applications [5,7]. Without adequate clock synchronisation, the drift of phone clocks (i.e. their measurement error) can reduce the modelling capability, especially when sensing physical events using multiple devices simultaneously. For instance, while earthquake sensing is possible using smartphone accelerometers [1], it is not clear whether data from multiple smartphones can be aggregated to accurately model the propagation of the earthquake waves or building collapses: this requires precise time synchronisation among all the smartphones.

In this paper we investigate the question: *What is the best clock synchronisation can we expect when sensing with a number of end-users' smartphones and publicly available NTP servers?* Our work presents a series of experiments and makes the following contributions:

- First, we quantify the drift on offline smartphones over long timescales. In the worst case, it can be more than 150ms per day. We then develop a rapid method for synchronising phone clocks within the range of 10-100ms. Compared to previous work [7],

our solution requires less synchronisation processing time and is technically easier to implement.

- Second, we evaluate our method in an experiment. We show that with only 10 NTP records the pairwise synchronisation error between smartphones achieves 27ms on average (measured by accelerometer), following a Gaussian-like distribution with maximum error 411ms.

Background Work

Research on mobile phone sensing has become an emerging paradigm of mobile computing. Major applications include transportation, environmental monitoring and healthcare [6]. Furthermore, time-critical applications of smartphones have been actively researched, such as transportation and environmental monitoring. Despite this effort, recent studies show that current mobile phone applications are plagued with time accuracy issues [5] because smartphone clocks cannot provide accurate time [7].

Accurate timestamps from smartphone clocks are required for personal sensing such as inertial localization, while they play a crucial role in crowdsensed data, such as earthquake monitoring [1].

Since cell towers do not provide public time synchronisation service, one practical approach to synchronise distributed devices is Global Positioning System (GPS). However, this has drawbacks: indoor and urban environments may degrade the signal; GPS receivers have high power consumption; and the use of GPS may conflict with location privacy.

Alternatively, the Precision Time Protocol (PTP) specified by IEEE 1588 standard [3] can also be used to synchronise clocks within a specific network. With the elimination of queuing delays, the time error of PTP is



Figure 1: Photograph from Experiment 2. The 3 devices are physically coupled and attached to the ceiling.

in the sub-microsecond range on local area networks. However, PTP is not practical for use with common mobile sensing tasks, where different devices may be connected via a variety of communication technologies (e.g., WiFi, 4G, 3G, GPRS) and across hybrid networks.

A practical alternative is NTP [10], which is a hierarchical architecture across a four-stratum network. Stratum 0, also known as reference clocks, is an array of devices generating highly precise periodic signals. The remaining strata are synchronised to their upper strata. Li et al. [7] designed an experiment to compare linear fitting and linear programming methods over different synchronising instances using NTP. They conclude that 6 hours is the best duration of a synchronising instance with a maximal error of 100ms.

Experiments

We conducted two experiments to quantify clock drift and synchronisation on smartphones. To avoid the effects from differences of operating system versions and manufacturers, both experiments used 3 identical Motorola Moto G mobile phones in the same Wi-Fi network. For data collection on the smartphones we used the AWARE Framework [2] with a custom NTP plugin, set to poll the NTP server every 1 minute.

Experiment 1 measured the drift of the smartphones over a period of 9 days. During this time, the devices were placed on a flat surface and our software recorded both the clock of the phone and its drift in relation to the NTP server. The phones were not set to update their time automatically, thus each device's drift accumulated during the experiment. The phones remained unused, to ensure that memory and CPU hogging did not affect our measurements.

Experiment 2 quantified the error in attempting to synchronise independent sensor streams using NTP. For

this experiment we decided to create a physical event that we could be certain was simultaneously detected by all phones. To achieve this, we physically coupled the 3 devices (Figure 1) and turned them into a pendulum. The 3 devices were hung from the ceiling, and they were swung in a pendulum fashion. The 3 devices were configured to record accelerometer data with a 20ms frequency, thus we could be certain that all phones physically experienced the same events (i.e., each swing) simultaneously. This offered us ground truth, since we are certain that the phones had to agree on the precise timestamp of each swing.

Results

Experiment 1

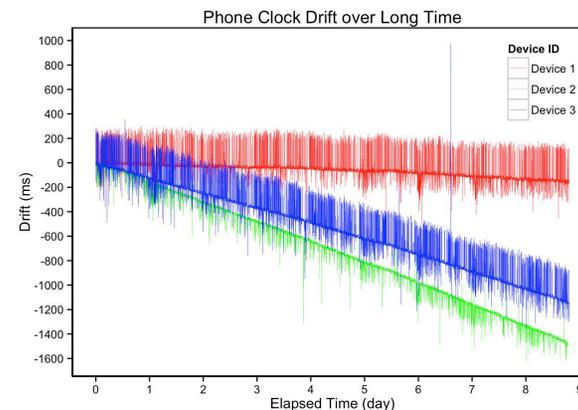


Figure 2: Clock drift error of three mobile phone clocks, quantified using NTP records (Experiment 1).

Figure 2 gives the results from Experiment 1, where the drift on each device is quantified using NTP records over 9 continuous days. We observe that the skewing rate for each device is linear, as has been previously

reported [7]. In the worst case, we observe a drift of approximately 1500ms over the 9-day period. Due to the inconsistent network latency, the NTP records for each device are often noisy, forming a fuzzy distribution around the linear drift. Most records range within 100ms of the true drift, while the maximum error was up to approximately 1800ms. This suggests that a single NTP record is not reliable because its timing can vary considerably.

Because we do not have ground truth on the true drift, we follow the recommendations from previous work [7] and model true drift as a regression. We use the Theil-Sen estimator [11] which is robust to outliers and our data has substantial outliers (Figure 2).

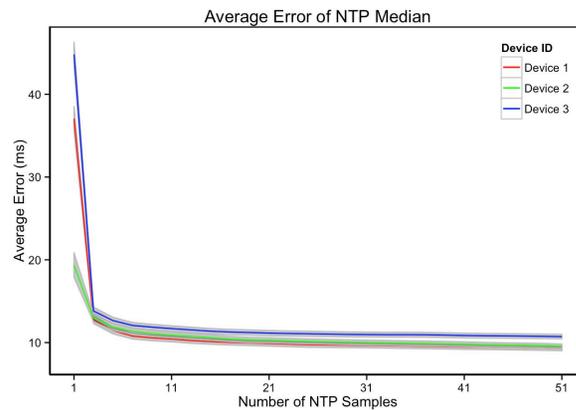


Figure 3: Average error when using the median of varying number of NTP samples to synchronise phone clocks. Samples are taken at 1 minute intervals (Experiment 1).

Against this ground truth model, we investigate alternatives to estimate drifts on-the-fly and with less NTP records. Specifically, our approach considers a window of T moments. At true time t in T , let the smartphone clock reading and the NTP record obtained

by this phone be $C(t)$ and $NTP(t)$. We compute the median of $C(t) - NTP(t)$ over a set of T moments. As the clock drift is negligible in a short period (several minutes or hours), we consider the calculated median as the true clock drift.

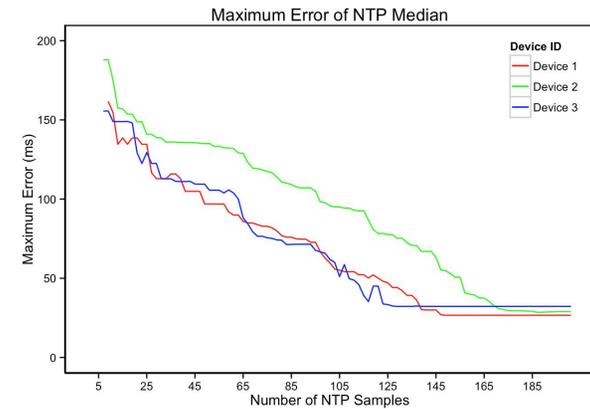


Figure 4: Maximum error calculation for the results of Figure 3 (Experiment 1).

Figure 3 shows the average error of our approach, which decreases significantly when the window of NTP samples increases from 1 to 5. Eventually, the error converges to less than 10ms for windows sized greater than 67. These results confirm that a single NTP record is rather unreliable, and that substantial precision gains can be achieved by using a window of 10 or more NTP samples. We also estimate the confidence intervals of t errors to be less than 1ms with at least 3 samples. Finally, we estimate the worst case by identifying the maximum error in our approach (Figure 4). Again, we observe that the maximum error declines as the window of NTP samples increases: it is less than 100ms when using 97 samples, and less than 40ms when using over 157 samples.

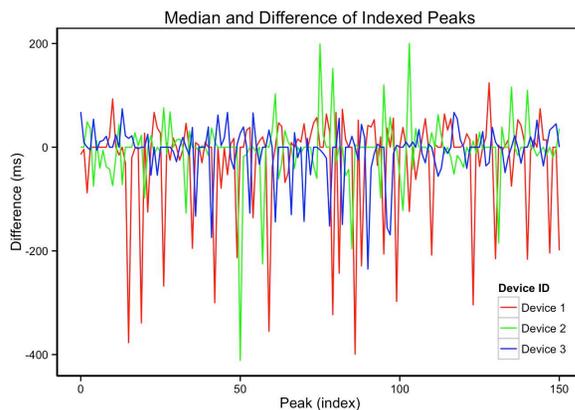


Figure 5: Median and difference of the 151 ground-truth events (Experiment 2).

Experiment 2

Based on the findings of Experiment 1, we adopt a window of 10 subsequent NTP records (each taken at 1-minute intervals) to synchronise the phone clocks. Then, by swinging our coupled smartphones in a pendulum fashion we observe 151 swings. Each swing is a physical event that acts as ground truth: each swing should be reported at the exact same time by each device. By inspecting the accelerometer data log from each smartphone, we index acceleration maxima (“peaks”) as ground-truth.

Figure 5 shows the 151 peaks, each with the relative pair-wise synchronisation difference between the 3 devices. The distribution of all these pairwise synchronisation differences is shown in Figure 6 (in yellow). In addition, Figure 6 shows the NTP offset for each phone from Experiment 1 (in red, green and blue). We observe that for each individual phone the offset forms a Gaussian-like distribution around the

accurate time. The synchronisation difference between devices in Experiment 2 again follows a Gaussian-like distribution with higher variance, mean of 27ms and maximum of 411ms.

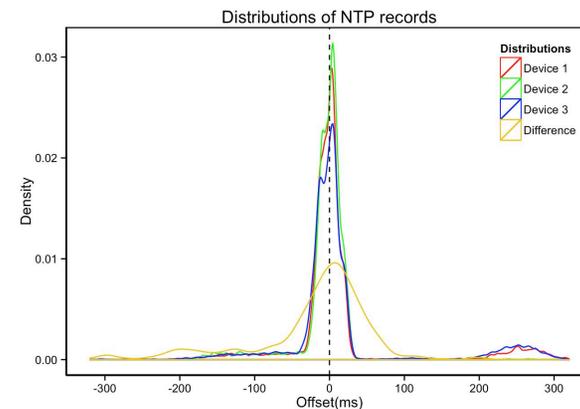


Figure 6: Distribution of NTP records for each device (Experiment 1) and their synchronisation difference (Experiment 2).

Discussion and Conclusion

We observe that the accumulated clock drift can be more than 150ms per day in the worst case. This result is greatly higher than previously reported (80ms) [7]. The theoretical error range of NTP synchronisation varies from several milliseconds to 100ms [9]. We find that the offset of NTP records over 9 days follows a Gaussian-like distribution, which confirms previous results [12]. Using a window of 10 NTP records with 1-minute intervals, our method can reach the error of less than 12ms on average. Previous work has required longer synchronisation processing time (6 hours in [7]) and more complex implementation with linear fitting or linear programming. Hence, our approach is practical for smartphone clock synchronisation scenarios.

We show that the synchronisation differences across devices appear a Gaussian distribution with an average of 27ms and maximum of 411ms. Since the offset of a single device appears a Gaussian distribution, the pairwise differences should remain stable between numerous devices. Mathematically, the difference of two independent variables from a Gaussian distribution follows another Gaussian distribution with a mean equal to the difference of variable means, and with a variance equal to the sum of variable variance.

We note that time synchronisation using GPS has a theoretical error up to 40ns [4]. However, even if GPS is available, NTP can be an alternative for privacy protection and battery conservation. Our results can guide researchers and developers in understanding what types of mobile phone sensing applications are possible for sensing physical phenomena.

Acknowledgements

This work is partially funded by SocialNUI, the Academy of Finland (Grants 276786-AWARE, 286386-CPDSS, 285459-iSCIENCE, 304925-CARE), the European Commission (Grant 6AIKA-A71143-AKAI), and Marie Skłodowska-Curie Actions (645706-GRAGE).

References

1. Linus Bengtsson, Xin Lu, Anna Thorson, Richard Garfield and Johan Von Schreeb. 2011. Improved response to disasters and outbreaks by tracking population movements with mobile phone network data: a post-earthquake geospatial study in Haiti. *PLoS medicine* 8, 8: 1128.
2. Denzil Ferreira, Vassilis Kostakos and Anind K. Dey. 2015. AWARE: mobile context instrumentation framework. *Frontiers in ICT* 2, 6: 1-9.
3. IEEE. 2008. IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems. *IEEE Std 1588-2008*: 1-269.
4. Elliott Kaplan and Christopher Hegarty. 2005. *Understanding GPS: principles and applications*. Artech house.
5. Yungeun Kim, Yohan Chon and Hojung Cha. 2012. Smartphone-based collaborative and autonomous radio fingerprinting. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on* 42, 1: 112-122.
6. Nicholas D. Lane, Emiliano Miluzzo, Hong Lu, Daniel Peebles, Tanzeem Choudhury and Andrew A. T. Campbell. 2010. A survey of mobile phone sensing. *Communications Magazine, IEEE* 48, 9: 140-150.
7. Dong Li and Pradeep Sinha. 2014. Rbtp: Low-power mobile discovery protocol through recursive binary time partitioning. *Mobile Computing, IEEE Transactions on* 13, 2: 263-273.
8. Artis Mednis, Girts Strazdins, Reinholds Zviedris, Georgijs Kanonirs and Leo Selavo. 2011. Real time pothole detection using android smartphones with accelerometers. In *Distributed Computing in Sensor Systems and Workshops (DCOSS), 2011 International Conference on*, 1-6.
9. David L. Mills. 2010. *Computer Network Time Synchronization: The Network Time Protocol on Earth and in Space, Second Edition*. CRC Press, Inc.
10. David Mills, Jim Martin, Jack Burbank and William Kasch. 2010. *Network time protocol version 4: Protocol and algorithms specification*.
11. Pranab K. Sen. 1968. Estimates of the regression coefficient based on Kendall's tau. *Journal of the American Statistical Association* 63, 324: 1379-1389.
12. Joseph D. Touch. 1992. *Mirage: A Model for Latency in Communication*. Ph.D Dissertation. University of Pennsylvania.