

Fingers of a Hand Oscillate Together: Phase Synchronisation of Tremor in Hover Touch Sensing

John Williamson
University of Glasgow
Glasgow
jhw@dcs.gla.ac.uk

ABSTRACT

When using non-contact finger tracking, fingers can be classified as to which hand they belong to by analysing the phase relation of physiological tremor. In this paper, we show how 3D capacitive sensors can pick up muscle tremor in fingers above a device. We develop a signal processing pipeline based on nonlinear phase synchronisation that can reliably group fingers to hands and experimentally validate our technique. This allows significant new gestural capabilities for 3D finger sensing without additional hardware.

KEYWORDS

capacitive; hover; tremor; phase synchronisation

ACM Reference Format:

John Williamson. 2016. Fingers of a Hand Oscillate Together: Phase Synchronisation of Tremor in Hover Touch Sensing. In . ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/2858036.2858235>

1 MOTIVATION

3D touch devices which can sense the fingers before they contact the screen have been researched for some time [4, 17, 18] and are becoming commercially available (e.g. Microchip's GestIC, Fogale's Sensation). Such 3D multitouch devices can be used by one or many people simultaneously; for example, in a two-player collaborative game. Reliably distinguishing interacting fingers from users interacting simultaneously can be tricky. Splitting the screen can provide multi-user input (e.g. one person gets the left half, the other gets the right half), but this is wasteful and doesn't allow for rich, collaborative multitouch interaction. These devices also cannot distinguish bimanual and single-handed gestures. New gestures which use knowledge of the generating hand extend the gestural repertoire of multi-touch screens.

We propose a method to group fingers according to the hand they belong to. Applications include: distinct unimanual and bimanual gesture input (e.g. one-handed pinch for microzoom functionality, large scale zooming with two-handed pinch); media browsing and sorting applications, where multiple users can slide and arrange items without gestures interfering with each other; and multiplayer




Figure 1: Frequency spectrum of the distance to surface of a finger hovering steadily 1cm above a touch screen. Tremor is visible around 9Hz.

above-surface games, with independent multitouch controls even in overlapping screen areas.

2 RELATED WORK

Tremor sensed via accelerometer has been used as a way of detecting whether mobile devices are being held by a living person, and as a proxy for grip pressure, by monitoring frequency changes in isometric tremor [15]. The phase dynamics of tremor are discussed in detail by Beuter et al [3]. Selker et al [14] use mobile accelerometer based tremor tracking to estimate sleep quality of users from long-term measurements. Phase synchronisation analysis has been used to detect conversational alignment between people speaking on the phone by measuring phase relations of gait cycles [10].

There has been extensive work in multi-user displays to separate interacting fingers from different users. DiamondTouch [5] used capacitive coupling through the body to separate hands (requiring users to sit in specific transmitter seats). Suto et al [16] use an optical system to estimate finger identity from user position. Medusa [1] used proximity sensors to detect the orientation of arms and assign touch points by projecting the arm orientation to the touch region. Harrison et al [8] use swept frequency capacitive sensing to distinguish users on a mobile device with additional hardware.

3 AIMS

All humans involuntarily produce strong tremor oscillations in their arms, in the 6-12Hz region [7], as a byproduct of muscular activation (Figure 1). There are various type of tremor present in healthy humans, including rest tremor, action tremor, postural tremor, kinetic tremor, isometric tremor and intention tremor; the

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(s).

CHI '16, May 07-12, 2016, San Jose, USA.

© 2016 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-3362-7/16/05.

<https://doi.org/10.1145/2858036.2858235>

most relevant tremor in the hovering finger case is *rest tremor*, which is produced while muscles are at rest. This has been used, for example, to determine grip strength using an accelerometer to track modulations of the tremor spectrum [15]. Muscle tremor effects are quite weak when measured with a finger resting on a surface, as on a conventional touchscreen, because the surface provides mechanical support and frictional damping to the finger. However, when fingers are hovering unsupported over the device, tremor is much more pronounced.

We propose to track the tremor component in each detected finger and identify common sources. We hypothesise that one user's fingers are subject to a large common tremor (from wrist and arm muscles), whereas fingers another user will be quite distinct. We can then group fingers according to their estimated common sources, and identify fingers belonging to different users or different hands of a single user.

4 SIGNAL PROCESSING

Our approach to estimating the common sources of tracked finger positions is to look for *phase synchronisation* between oscillations in the tremor range (6-12Hz). Where there is strong synchronisation between two tracked fingers, we might reasonably expect that they are attached to the same limb. Physiological evidence suggests there is little synchronisation between left and right side limbs (excluding small ballistocardiac effects) as there is hypothesised to be a pair of lateralised tremor generators in the body [7].

To estimate phase synchronisation, we use a Hilbert transform to approximate the analytic signal [6], and then compare the resulting phase angles. This method is robust and insensitive to amplitude modulations.

4.1 Pre-processing

For each finger, we capture the tracked x, y, z finger positions from the touch pad, and the timestamp t . We reinterpolate all of the position time series to a constant sampling rate of 120Hz. We then estimate the derivative of each of these values by finite differences. These signals are then bandpass filtered, as the Hilbert transform used later relies on the assumption of a narrowband signal. We use linear phase FIR filter to avoid phase distortion, with 91 taps, using Blackman-Harris windowed filter design. The -3dB band cutoffs are set to 6Hz and 30Hz, to capture the tremor oscillations typically centred around 10Hz. This produces a bandpass filtered signal $x_f(t)$ for each axis.

4.1.1 Hilbert Transform. From the bandpass filtered signal $x_f(t)$, we compute the Hilbert transform $x_h(t) = H(x_f(t))$, to get the complex analytic signal. In the offline case, we compute the true (discrete) Hilbert transform via the FFT of the entire signal. A Hilbert transform approximation FIR filter [13] can be used in an online implementation (the bandpass and Hilbert filters can be convolved into a single FIR filter in this case). From the analytic signal, we can obtain the phase angle $\phi(t) = \tan^{-1}(\Im(x_h(t)), \Re(x_h(t)))$ which is phase unwrapped to produce $\phi_u(t)$. The Hilbert transform is essential in achieving amplitude invariant performance, which is critical for the hover sensing application as the sensor's distance mapping is non linear. This means tremor with a distant finger has




Figure 2: Block diagram of the signal processing steps in the tremor synchronisation tracking.

very different amplitude distribution than with the finger close to the screen.

4.1.2 Phase Locking Value. Computing this value for a pair of fingers, we obtain $\phi_u^1(t), \phi_u^2(t)$. We can compute the *phase locking value* (PLV) [2] as:

$$P = \frac{1}{T} \left| \sum_{t=0}^T e^{i(\phi_u^1(t) - \phi_u^2(t))} \right|.$$

This is a measure of the difference of phase of $\phi_u^1(t), \phi_u^2(t)$ over time and is widely used in the EEG literature when investigating neural synchronisation [9].

The PLV value measures the synchronisation of two signals. However, if the two signals are for example constant (e.g. finger position is completely still or absent), the PLV value will be high and it can be difficult to find reliable thresholds for the PLV to indicate synchronisation. A robust approach to identifying synchronisation is *data surrogacy* [11, 12] which compares the phase synchronisation between the original signals and replicates which have had temporal structure destroyed. This separates the time varying synchronisation effect from overall signal statistics. This involves computing a set of *surrogates*. Each surrogate is computed by randomly shuffling the order of the time-series, and then re-computing $\phi_u^*(t)$, the unwrapped phase for the shuffled data, and finally the PLV value P_j^* for each surrogate j . By generating multiple surrogates we can derive a robust statistical measure of how informative the PLV is.

We estimate the mean μ_{P^*} and standard deviation σ_{P^*} of the PLV for each surrogate replication, and compute the z-score of the real (original) PLV as $z_P = \frac{P - \mu_{P^*}}{\sigma_{P^*}}$. This is a robust measure for the level of true synchronisation. We would expect signals with strong synchronisation to have z-score $>> 1$, indicating that the true PLV is well outside the distribution of the surrogate PLVs. The overall process is summarised in Figure 2. In practice, a small number of surrogates (2-3) suffice, with minor gains in robustness beyond this level.

4.2 Real time implementation

The above process is used in the experimental work described below. To demonstrate that this technique is viable in real-time processing of touch input, a simplified online algorithm was also implemented. This implementation uses a truncated FIR approximation to the Hilbert transform (31 taps, designed by the Remez exchange algorithm) and a shorter FIR bandpass filter (also 31 taps), with a bandpass region of 8Hz-24Hz.



Figure 3: Left: The 3D sensing device, which also has a built-in display. Right: Raw output of the 10x16 sensor with two fingers hovering at approximately 0.5cm.

The PLV P_t is accumulated in a leaky integrator, and a connection between fingers is triggered when this value exceeds a threshold. To eliminate spurious phase locking, a single surrogate is estimated using a bandpass filter with randomly permuted elements, to produce P_t^* . To be considered a lock, $P_t > P_t^* + \Delta P$ for some threshold ΔP .

5 EXPERIMENT

N=8 subjects (2 female, 6 male, mean age 31.3) were recruited locally to test the tremor tracking system. We tested specifically whether the system could distinguish the left hand from the right hand of a single participant. Pilot testing indicated that separating different users' finger grouping was very reliable, so we concentrated on capturing intra-user hand grouping (synthetic comparisons across users can be made post hoc by comparing time series from different capture sessions).

Participants were seated with the sensor resting on a table about 30cm in front of them. The experimental trials consisted of a series of hand poses which were given via instruction images. Each pose was held for 20 seconds and repeated three times. There were a mix of static poses (i.e. hands held at rest) and dynamic poses (fingers moving). Participants were instructed to keep their fingers around 1cm above from the screen and avoid touching it. A simple bar chart of detected finger height was shown onscreen with a "good distance" zone, and participants were asked to use the chart to keep their fingers in this zone.

There were a total of four poses. Each pose had a certain number of fingers active on each of the left and right hand and was either static or was dynamic with circling motions (while keeping the fingers over the sensor area). The dynamic poses were included to demonstrate that the tremor synchronisation can be detected with moving fingers. Each pose has a unique code "R-LT" (right fingers, left fingers, type). For example, one finger from each hand, static was coded 1-1S, and two fingers from the right hand, dynamic was coded 2-0D. The poses used were: 2-0S, 2-0D, 1-1S and 1-1D.

5.1 Hardware + logging

We use a touch sensing device with 160 long range electrical field sensors arranged in a 10x16 grid (Figure 3). This sensor produces intensity values which vary as conductive bodies approach the surface. The sensor samples at 120Hz, but the capture software provides data at 100Hz. We experimentally estimated the signal to noise ratio (SNR) to be around 48dB. The sensor has a maximum




Figure 4: Box plot showing the z scores for the two separate finger case (1-1S and 1-1D) and two fingers on the same hand (2-0S and 2-0D), z-axis (finger height) only. S and D denote static (finger held steady) and dynamic (finger moving in circles) respectively. (bar=interquartile range, whiskers=Q1-1.5IQR, Q3+1.5IQR, median at red line)

sensing range of approximately 5cm above the surface, but is most effective in the 0-2cm range. Onboard firmware computes up to five independent finger x, y, z positions, which is streamed along with the raw sensor values to a remote logging application. A Python application coordinated experimental conditions and data logging. We have also tested with a second sensor device of lower sampling rate (60Hz) resolution (13x8) and obtained very similar performance to that presented here.

6 ANALYSIS AND RESULTS

Captured signals were analysed with the process outlined above, using $k = 12$ surrogates for z-score estimation. Figure 6 shows the phase angle of the z axis signals and a synchrogram [12] from a typical two finger case (1-1S versus 2-0S). The synchronisation of the signals from the same hand is very marked in both the synchrograms and the PLV score distributions. The full results from the trial for the two finger case is given in Figure 4. The synchronisation detection is very reliable. A Wilcoxon rank-sum test shows a statistically significant difference between the z scores for 1-1S and 2-0S ($p < 0.002$) and between 1-1D and 2-0D ($p < 0.0002$).

The results suggest that the z axis movement (height above surface) is a better measure than the x and y movement (translation in the plane parallel to device surface), although synchronisation is still visible in those axes (Figure 5). The mechanism underlying this axis difference remains to be investigated but we hypothesise it may be due to the biomechanics of the arm when held in the hovering pose. In the separate finger cases z-scores are generally still above 1, although much lower than the z-scores for the connected finger cases. This is due to glitches in sampling rate being propagated across all tracked fingers; finger grouping is still easily




Figure 5: Box plot showing the z scores for the two separate finger case (1-1S and 1-1D) and two fingers on the same hand (2-0S and 2-0D), for x , y , z positions (z highlighted in green). x , y show less strong synchronisation than z (finger height).

distinguishable as synchronised fingers have z-scores $\approx 5\times$ that of non-synchronised fingers. Combining multiple axes and tracking the full 3D tremor pattern may improve performance, although this requires tracking continuously varying phase offsets between the axes.

We subsequently captured data for 2-2S, 3-1S, 4-0S and 2-1S from all participants. The analysis of the connectivity graphs of the fingers is significantly complicated by the touch API's unstable numbering of the touch points and the results cannot be presented in detail. In summary, the results are positive; tremor grouping remains robust with three and four fingers, as long as the touch sensor is able to reliably distinguish fingers. We also informally tested the effect of gesture velocity and found that the effect of voluntary motion to be insignificant even for rapid movements (e.g. quick swiping of fingers back and forth), although the small size of the sensor makes it difficult to sustain quick movement without losing tracking.

7 DISCUSSION

Detection of synchronisation is very reliable. PLV based scoring works very well, but simpler approaches, such as simple phase comparison (which computes how closely clustered points are around the $x = y$ line on a synchrogram) would be likely to be effective. Spectral contamination from voluntary movements is implausible and thus the detection is generally robust; however sources of vibration in the 6-12Hz range (e.g. in vehicles) has not been tested and may reduce the sensitivity of the tremor grouping.

Our approach has both processing and inherent latency before synchronisation can be detected. The FIR filters have excellent phase linearity but introduce substantial processing latency; the real-time implementation discussed has group delay of around 260ms. IIR approaches could mitigate this. In terms of inherent latency,

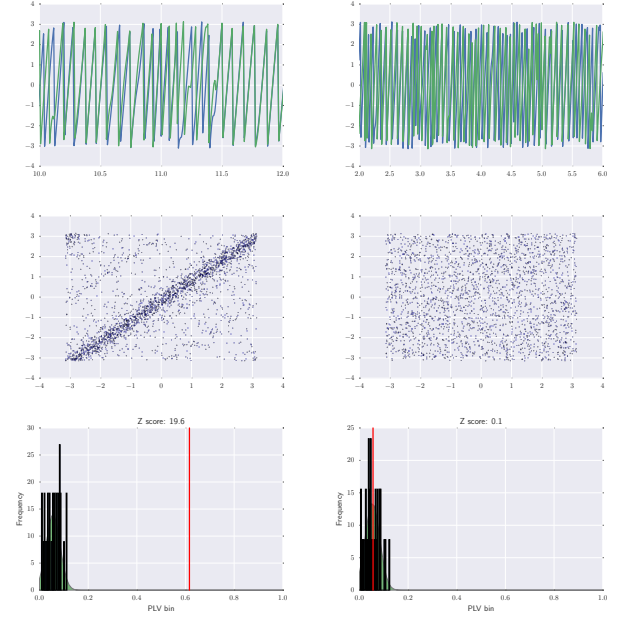


Figure 6: Time series from 30 seconds of typical two finger data. Left: two fingers from one hand, right: two fingers from different hands. Top: 2 second extract of phase angle of the signal ϕ_z for the two series overlaid. Middle: synchrogram (ϕ_{z1} versus ϕ_{z2}). Bottom: True P value (red) and the distribution of 100 P^* surrogate values, along with the corresponding z-score. The synchronisation is very apparent in the strong clustering of the synchrogram plots around $x = y$, closely overlapping phase time series and PLV well outside of the distribution of surrogate PLVs.

the Hilbert transform approach needs at least two or three complete cycles to reliably identify synchronisation. At 10Hz tremor, this means 200-300ms of data must be observed, so it is unsuitable for very fast swipes above the device. Other approaches to phase synchronisation detection, such as phase-locked loop approaches, may offer more sensitivity and lower-latency lock on. However, for many tasks (e.g. multi-user media browsing) the grouping structure is relatively stable and sub-second updates are not required. The analysis approach presented here works on tracked finger positions. The phase synchronisation can also be detected *directly* on the sensor matrix, before finger extraction has been completed by computing PLV between each sensor matrix value and performing spectral clustering on the resulting similarity matrix. This may offer benefits in separating closely spaced fingers.

8 CONCLUSION

Human tremor is a prevalent and easily detectable signal that has significant applications in improving user interface design. 3D touch tracking has to cope with the “noise” of tremor in positioning, but instead of simply filtering this it out, it can be harnessed to augment input without any additional hardware. Our processing pipeline

is a robust and easily-implemented way to group fingers to hands and opens up an array of hand-origin sensitive gestural controls.

ACKNOWLEDGMENTS

We thank Andrew Ramsay for assistance in software development. All data from this study is available at: <http://dx.doi.org/10.5525/gla.researchdata.246>

REFERENCES

- [1] Michelle Annett, Tovi Grossman, Daniel Wigdor, and George Fitzmaurice. 2011. Medusa: a proximity-aware multi-touch tabletop. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*. ACM, 337–346.
- [2] Sergul Aydore, Dimitrios Pantazis, and Richard M Leahy. 2013. A note on the phase locking value and its properties. *Neuroimage* 74 (2013), 231–244.
- [3] Anne Beuter, Leon Glass, Michael C Mackey, and Michele S Titcombe. 2013. *Nonlinear dynamics in physiology and medicine*. Vol. 25. Springer Science & Business Media.
- [4] Sangwon Choi, Jaehyun Han, Sunjun Kim, Seongkook Heo, and Geehyuk Lee. 2011. ThickPad: a hover-tracking touchpad for a laptop. In *Proceedings of the 24th annual ACM symposium adjunct on User interface software and technology*. ACM, 15–16.
- [5] Paul Dietz and Darren Leigh. 2001. DiamondTouch: a multi-user touch technology. In *Proceedings of the 14th annual ACM symposium on User interface software and technology*. ACM, 219–226.
- [6] Dennis Gabor. 1946. Theory of communication. Part 1: The analysis of information. *Journal of the Institution of Electrical Engineers-Part III: Radio and Communication Engineering* 93, 26 (1946), 429–441.
- [7] Mark Hallett. 1998. Overview of human tremor physiology. *Movement Disorders* 13, S3 (1998), 43–48.
- [8] Chris Harrison, Munechiko Sato, and Ivan Poupyrev. 2012. Capacitive fingerprinting: exploring user differentiation by sensing electrical properties of the human body. In *Proceedings of the 25th annual ACM symposium on User interface software and technology*. ACM, 537–544.
- [9] Jean-Philippe Lachaux, Eugenio Rodriguez, Jacques Martinerie, Francisco J Varela, and others. 1999. Measuring phase synchrony in brain signals. *Human brain mapping* 8, 4 (1999), 194–208.
- [10] Roderick Murray-Smith, Andrew Ramsay, Simon Garrod, Melissa Jackson, and Bojan Musizza. 2007. Gait alignment in mobile phone conversations. In *Proceedings of the 9th international conference on Human computer interaction with mobile devices and services*. ACM, 214–221.
- [11] Milan Palus and Dirk Hoyer. 1998. Detecting nonlinearity and phase synchronization with surrogate data. *Engineering in Medicine and Biology Magazine, IEEE* 17, 6 (1998), 40–45.
- [12] Arkady Pikovsky, Michael Rosenblum, and Jürgen Kurths. 2003. *Synchronization: a universal concept in nonlinear sciences*. Vol. 12. Cambridge university press.
- [13] David Ernesto Troncoso Romero and Gordana Jovanovic Dolecek. 2012. *Digital FIR Hilbert Transformers: Fundamentals and Efficient Design Methods*. INTECH Open Access Publisher.
- [14] Ted Selker, Patricia Collins, and Will Dayton. 2011. Psychosocial indicators via hand tremor. In *Human-Computer Interaction—INTERACT 2011*. Springer, 596–599.
- [15] Steven Strachan and Roderick Murray-Smith. 2004. Muscle tremor as an input mechanism. In *Proc. UIST*, Vol. 2. 482.
- [16] Shota Suto and Susumu Shibusawa. 2013. A Tabletop System Using Infrared Image Recognition for Multi-user Identification. In *Human-Computer Interaction – INTERACT 2013. Lecture Notes in Computer Science*, Vol. 8118. 55–62.
- [17] Yoshiaki Takeoka, Takashi Miyaki, and Jun Rekimoto. 2010. Z-touch: an infrastructure for 3d gesture interaction in the proximity of tabletop surfaces. In *ACM International Conference on Interactive Tabletops and Surfaces*. ACM, 91–94.
- [18] Andrew D Wilson. 2010. Using a depth camera as a touch sensor. In *ACM international conference on interactive tabletops and surfaces*. ACM, 69–72.