Corona: Positioning Adjacent Device with Asymmetric Bluetooth Low Energy RSSI Distributions

Haojian Jin
Yahoo! Labs
701 First Ave
Sunnyvale, CA 94089
haojian@yahoo-inc.com

Cheng Xu
Capital One
201 3rd St
San Francisco, CA 94107
chengx.cn@gmail.com

Kent Lyons
Technicolor Research
175 S. San Antonio Rd STE 200
Los Altos, CA 94022
kent.lyons@technicolor.com

ABSTRACT

We introduce Corona, a novel spatial sensing technique that implicitly locates adjacent mobile devices in the same plane by examining *asymmetric Bluetooth Low Energy RSSI distributions*. The underlying phenomenon is that the off-center BLE antenna and asymmetric radio frequency topology create a characteristic Bluetooth RSSI distribution around the device. By comparing the real-time RSSI readings against a RSSI distribution model, each device can derive the relative position of the other adjacent device. Our experiments using an iPhone and iPad Mini show that Corona yields position estimation at 50% accuracy within a 2cm range, or 85% for the best two candidates. We developed an application to combine Corona with accelerometer readings to mitigate ambiguity and enable cross-device interactions on adjacent devices.

Author Keywords

Mobile phones; sensing; BLE; RSSI; multi-device.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous

INTRODUCTION

As Weiser predicted in the 90's [12], we are now surrounded by multiple personal devices (e.g. mobile phones and tablets). Interacting with multiple devices at the same time and switching among them has become a common task in daily life. While there are mature communication technologies for these devices (e.g. Bluetooth, Wi-Fi), the devices remain independent because of the lack of a spatial awareness even when they are just inches away [1].

In this work, we introduce *Corona*, a technique that implicitly locates the position of adjacent mobile devices placed in the same plane. Many commercial devices place the wireless antenna at the corner due to the large size of the battery and have an asymmetric radio frequency field and associated RSSI distribution around the device (**Figure 1**). By measuring a real-time Bluetooth Low Energy (BLE) signal strength sequence and comparing it to pre-collected distributions of RSSI values, Corona is able to determine the relative positions of adjacent devices. Corona runs between two peers without any infrastructure. It only requires the compass, gyroscope, accelerometer and BLE, which are built into almost all modern mobile devices.

The contributions of this work are: 1) A study of the asymmetric radio frequency field pattern around commodity mobile devices. 2) The development of a Bayesian method to estimate relative positions through the real-time RSSI readings and our evaluation of this method. 3) A proof of concept application showing how the implicit calibration can alleviate the multi-solutions ambiguity.

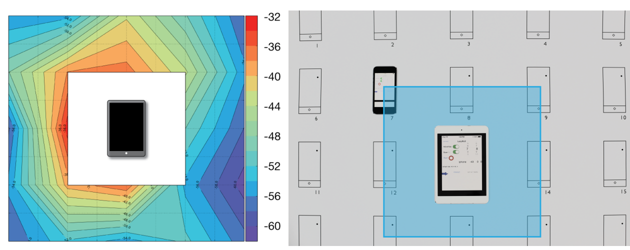

Figure 1: Left: Example BLE RSSI distribution map. Right: Positions we collected readings from. No interpolation is performed within the blue square resulting in the white space on the left.

RELATED WORK

One common method for radio localization is fingerprinting [7, 9]. Corona is similar in spirit to fingerprint methods but differs in some key ways. Fingerprinting creates a probability map of a given area based on signal strength of *several* access points or base stations. To determine a location, new RSSI readings of various access points are collected and these systems search for a position with the most similar signal strength readings. The accuracies of fingerprint methods are usually in meters [7, 9]. Corona also uses pre-recorded RSSI readings. However, it is designed for ad-hoc cross-device interactions. Instead of using radio signals from the environment with several sources, we use the single transmission signal of the other mobile device. Our system maps the signal strength sequence from that device to the pre-collected model that represents the asymmetric dis-

UIST '15, November 08–11, 2015, Charlotte, NC, USA

ACM 978-1-4503-3779-3/15/11…$15.00
DOI: http://dx.doi.org/10.1145/2807442.2807485

tribution of RSSI values at different locations around the mobile device.

Besides radio localization, past research has provided multiple solutions to enable spatial awareness of mobile devices. GroupTogether [5] uses multiple radio modules to infer coarse-grain device layout and locates mobile devices through tracking users with a pair of overhead Kinect depth cameras. Similarly, HuddleLamp [10] uses computer vision to track the mobile devices through an integrated RGB-D camera. These approaches require special hardware or non-pervasive infrastructure. Corona relies solely on commodity sensors that are available on most mobile devices.

Like Corona, other research has also explored cross-device interaction with built in sensors. BeepBeep [6] and SwordFight [13] developed an ad-hoc phone-to-phone distance measurement system using microphones and speakers. Stitching [3] allows users to combine multiple displays into a cross-device workspace by dragging a pen across them. The required stitching gesture was used to determine the spatial configuration of the devices. Similarly, SurfaceLink [2] and AirLink [1] associate the devices through an explicit gesture on the surface and in air. These approaches, however, requires explicit user actions or explicit device operations (sound/vibration). Corona infers the relative location of two devices that have been placed adjacent to each other.

CORONA

Asymmetric RSSI Distributions

Corona leverages asymmetric radio frequency field patterns to locate adjacent devices. The asymmetry comes from the off-center placement of the antenna, the antenna design and the different material and shielding properties of the device. With Corona, we map this asymmetry, and in particular develop RSSI distribution models that correspond to different positions around the device. Figure 2 shows the RSSI distribution at three different adjacent positions around an iPad mini as measured from an iPhone. While the distance between these three positions is only 3cm, the distribution differences are still noticeable.

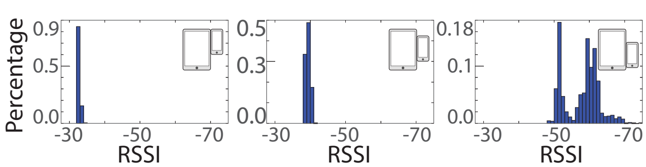

Figure 2: The RSSI distribution collected at the positions along the left sides with a 3cm interval. There is a visible difference between these positions.

Modeling RSSI distributions

RSSI distributions are known to be hard to model because of the random left-skewed nature of the distribution [4]. For our close proximity case, we find the RSSI readings are repetitive, including the left-skewed part. We therefore use a discrete RSSI probability distribution to represent the RSSI characteristic at different positions.

Given a specific position configuration at c,

$$P(x|c) = {count(x)}/{totalcount}$$

where $P(x|c)$ denotes the probability that specific RSSI value x happens at position c. $count(x)$ denotes the count of RSSI x in the data collection at position c, $totalcount$ is the total count of the RSSI samples recorded at that position.

Ambiguity

In practice, we find this phenomenon suffers from ambiguity. As shown in Figure 1, some positions have similar signal strength. To overcome this ambiguity for Corona, we focus on locations where the mobile devices are in immediate proximity: the devices are in the same plane and physically touching (**Figure 3**).

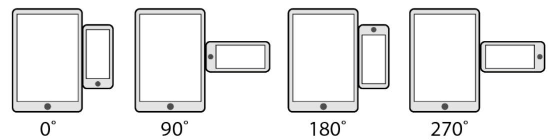

Figure 3: Mobile devices are placed in adjacent positions in one of four orientations.

System and Implementation

We selected an iPad Mini and an iPhone 5 to test our system. Corona runs as a peripheral manager (radio broadcasting) and a central manager (radio scanning) simultaneously through the standard iOS 8.1 API. Both managers run with default setting with the *allow duplicate* option so Corona obtains the raw RSSI data before system filtering.

The central manager reads RSSI values at more than 45 Hz between the iPhone and iPad. We aggregate the data from the two devices resulting in an overall RSSI sampling rate of over 90 Hz. On iOS, the BLE central manager enters into power saving mode after several minutes of scanning. To obtain continuous scanning, we restart the scanning thread at approximately 5 minute intervals.

Corona only runs while the devices are physically stationary. We detect this condition with a simple threshold filter on the accelerometer data. Once Corona detects a movement has finished, Corona starts the RSSI collection process and exchanges orientation data obtained from the magnetometer through BLE.

RSSI Distribution Model

As mentioned earlier, Corona needs to have a prior model of the RSSI distributions around the devices. To obtain this information, we collected RSSI data at pre-calibrated positions. We selected 16 positions along the four edges of the iPad Mini (5 on each width side and 3 on each height side) (**Figure 4**). The spacing between any two points on the same side is evenly distributed. For each position, the iPhone is placed adjacent to the iPad four times, each time at a different orientation and therefore with a different edge touching.

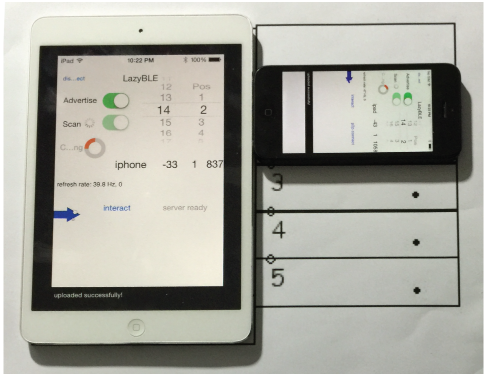

Figure 4: Ground truth sheet for data collection.

We printed the positions on paper in full scale (**Figure 4**) and collect 30 seconds of data at each designated configuration. This results in approximately 2700 (30s * 90Hz) RSSI readings for each configuration.

Bayesian Classification

At run time, Corona processes the signal in segments of k frames from the iPhone and iPad, and runs a Bayesian classification over each segment. This process yields a probability score for each pre-calibrated position. We empirically set k to 40, so we get a position update approximately once a second.

As mentioned above, we collect the RSSI distributions for each of the four orientations of the iPhone relative to the iPad. The relative orientation is sensed with the magnetometer and gyroscope. As such, Corona maintains and uses four separate models, one for each orientation. The estimation process can be formulated as follows. Let $C = \{c_1, c_2, \dots, c_n\}$ be the n (n=16 in our implementation) calibrated positions around the device. Given an RSSI sequence from the iPhone $X = \{x_1, x_2, \dots, x_k\}$ and an RSSI sequence from the iPad $Y = \{y_1, y_2, \dots, y_k\}$, the conditional probability that [X, Y] is recorded at c is $P(c|X,Y)$. Determining the recording position is equivalent to finding c^* that maximizes $P(c|X,Y)$:

$$c^* = \operatorname*{argmax}_c P(c|X,Y)$$

From Bayes' rule, $P(c|X,Y)$ can be formulated as:

$$P(c|X,Y) = \frac{P(X,Y|c)P(c)}{P(X,Y)}$$

where $P(c)$ represents the prior probability of selecting c without the observation of the RSSI reading, $P(X,Y|c)$ is the likelihood function, which expresses how probable the two sequences can be seen at position c, and $P(X,Y)$ is the normalization constant.

Since $P(X,Y)$ is a constant, we have:

$$c^* = \operatorname*{argmax}_c P(c|X,Y) = \operatorname*{argmax}_c P(X,Y|c)P(c)$$

We treat each RSSI reading in the sequence [X, Y] as independent and therefore the likelihood function is interpreted as the product of each RSSI x_m, y_m happening at position c:

$$P(X,Y|c) \sim \log\left(\prod_{m=0}^{k} P(x_m|c) * \prod_{m=0}^{k} P(y_m|c)\right)$$

and where we use the logarithm to convert the score to a linear distribution. $P(x_m|c)$, $P(y_m|c)$ denotes the probability that specific RSSI value happens at position c.

When Corona starts, $P(c)$ is evenly distributed initially because there is no prior information. In a later section, we introduce a method for inferring $P(c)$ through accelerometer data and help alleviate ambiguity.

Position Interpolation

In a real world application, the device is not necessarily placed at the pre-calibrated positions; we further refine our solution through a linear interpolation of adjacent candidates with top scores.

$$c_{finali} = \frac{c^* * P(X|c^*) + c_{next} * P(X|c_{next})}{P(X|c^*) + P(X|c_{next})}$$

where c^* is the position with the highest score, c_{next} is the adjacent point with the next highest score.

EVALUATION AND RESULTS

To evaluate how accurate Corona estimates the relative location of the other device, we conduct two separate experiments in a laboratory environment. The first evaluates the performance of positions captured in our RSSI models and the second tests the performance of interpolations on unaligned positions.

Evaluation with Aligned Positions

We collect 30 seconds of data at each designated position and split the data into chunks of 40 frames. We run the Bayesian classification described above over both iPhone/iPad data and compare the prediction with the ground truth. The data set used to build our model was collected on a wooden table in a home setting and the test set was collected in conference room in an office building.

Figure 5 shows the confusion matrices for our 16 positions across the four relative orientations. In all four confusion matrices, the strong diagonal shows a strong correlation between different positions and our inference based on the RSSI distributions. Specifically, orientation 270° has the best performance among the 4 relative orientations. For orientation 270°, there is some confusion between Position 1 (left-bottom corner) and Position 14 (bottom-right corner). As a result, 12% and 29% of position 1 and position 14 respectively are confused by each other. More generally, the corner positions (1, 6, 8, 14) tend to have lower accuracies than the other locations which seems to be due to an overlap in signal strength. Corona's second highest location estimate correctly predicts the ground truth for 97% of the misclassified instances.

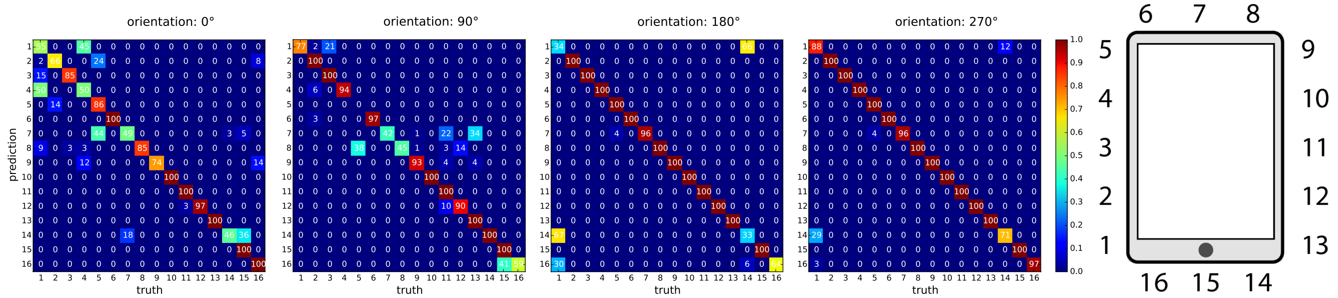

Figure 5: Confusion matrices of the four different relative orientations. The annotations inside the matrix show the percentage. We encode the position around the iPad in a clock-wise circular way, which is visualized in the right figure

Prediction with Unaligned Positions

To evaluate a more realistic use case, we expand our evaluation to unaligned positions around the device. We use another full scale reference sheet, which has randomly selected positions on the four edges of the iPad. This sheet contains 8 positions for each Left/Right side and 3 positions for each Top/Bottom side, so there are (8*2+3*2)*4 = 88 configurations in total. Similar to the earlier evaluation, we collected 30 seconds of data for each configuration and analyze it in the chunks of 40 frames windows. Again, the data set used for our model was from the in-home setting while the test set was collected in a conference room.

Orientations	Top-1	Top-2
0°	63.75%	90.14%
90°	41.79%	79.48%
180°	50.87%	88.84%
270°	50.92%	85.32%

Table 1. Comparison of accuracy rates for 4 orientations

As with our previous evaluation, a significant portion of the misclassifications are because of an overlap in signal strength. We also include the prediction with the 2nd highest score in the interpolation evaluation. We treat all of the predictions produced with an error smaller than 2cm as correct (2cm is the average half distance between adjacent pre-calibrated points). Table **1** reports the accuracies of the predictions in different orientations. The Top-1 column reports the success rate that the position with the highest score matches ground truth, while Top-2 column expands the candidate pool to the top 2 guesses.

As shown in Table **1**, Corona locates the device accurately in more than half of the estimates and locates the device at over 85% accuracy when considering the top two position candidates. For the remaining 15% of errorful predictions mainly result from the non-continuous nature of the RSSI distributions. For example, for a position between 2 and 3 in the 90° orientation, Corona correctly found the point with the most similar distribution as position 2. However, the score of position 1 is higher than that of position at 3 in around 70% estimates. In these cases Corona predicts a position between 1 and 2 thus producing an error larger than 2 cm.

SENSOR FUSION IN APPLICATION

In the evaluation, we find the most ambiguities happen at the corners of the iPad device (annotated in the same color in Figure 6a). Usually, one of the corners is on the long side and the other is on the short side. To mitigate this ambiguity, we developed a refinement that determines if the adjacent side is the long or short side through use of the onboard accelerometer and user movement of the devices.

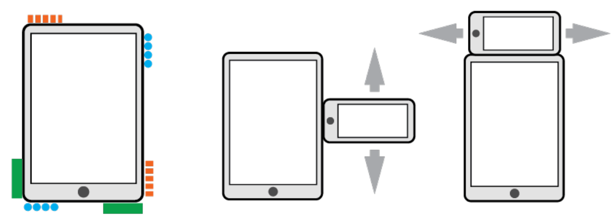

Figure 6: a). Ambiguity area for iPad Mini and iPhone b, c). The relative orientations further limit the motion directions.

Given the orientation configuration for two adjacent devices, knowing the iPhone movement direction can let us infer if the iPhone is on the long side or short side. Taking orientation 90° as an example, if the iPhone moves along its own left/right direction (**Figure 6**b), the iPhone can only be on the long sides of iPad. Similarly, if the iPhone moves along its own top/down direction (**Figure 6**c), the iPhone can only be at the short sides of iPad. The refinement is not required for Corona localization. However, if this happens implicitly by a user's movement, Corona detects the movement direction and updates the $P(c)$ in the Bayesian model.

DISCUSSION

Corona is a novel sensing technique based on common onboard sensors. It can be complementary to existing techniques [1, 2, 3] and further improve overall performance. It is also more computationally efficient than the common audio-based approaches [1, 2, 13], since it only requires processing 90hz sensor readings with linear complexity algorithm.

While running a statistical evaluation to evaluate the generalization of this approach is beyond the scope of this paper, we include some of our observations during Corona's development. We tried Corona with different devices including two iPhone 5's and two iPhone 6 pluses with one iPad Mini at multiple relative positions. The RSSI distributions at the same positions for the same types of devices appeared roughly same and generally the mean differences between different devices is under a 0.5 RSSI value. For the iPhone 6 plus combination, the RSSI distributions change significantly within a small amount of movement. While the effect of ambiguity is still unknown without a formal evaluation, the asymmetrical BLE phenomenon is still present and is an area to explore in future work.

We also tried more than two devices. In our informal observations, changing the number of devices did not seem to have a strong influence on the RSSI distributions. From our informal testing, most RSSI mean changes are again smaller than 0.5.

The interference from additional Wi-Fi/BLE devices in the environment seemed negligible. It is possibly this is because the main RSSI range for Corona is [-50, -20], which is very strong in terms of RSSI. Putting the device 10 cm away, the RSSI would drop to -60 or even less.

LIMITATIONS AND FUTURE WORK

Corona's interpolation of adjacent points assumes that the RSSI distribution in the area of interest is continuous. However, this assumption may not be true [5]. Modeling more positions along the edges with shorter windows may better model this non-continuous distribution.

Corona disables Wi-Fi during the scanning and broadcasting process. The interference from enabling Wi-Fi mainly results in more left-skewed components of the RSSI distribution. A longer data collection or accessing potential network APIs to exclude the Wi-Fi usage moments might alleviate this effect. Better Wi-Fi/Bluetooth coexistence radio technology might help in the future as well.

Both the evaluation and training were conducted in locations without large metallic objects around or in the table. We found RF-reflective surfaces produces an unpredictable offset to the whole distribution. One possible solution is developing a more sophisticated models based on RSSI distribution patterns rather than the absolute values. Another is integrating previous work [3] as a quick calibration method to determine the offset.

Currently, we constrain the interaction space to adjacent devices on a flat surface to minimize the ambiguity. In the future, Corona could be applicable in a broader space, e.g. co-planar in the air, or 3D space. Corona can also potentially expand to multiple devices, when combined with techniques like triangulation.

CONCLUSION

In this paper, we described how the asymmetry of BLE RSSI distributions is used to determine relative locations of two proximate devices. We implemented a robust statistical analysis algorithm that can locate the adjacent device at the top two estimates with an 85% accuracy. We implemented the system on an iPad and an iPhone; however, this solution could run on any commodity devices that are BLE enabled.

REFERENCES

1. Chen, K. Y., Ashbrook, D., Goel, M., Lee, S. H., & Patel, S. AirLink: sharing files between multiple devices using in-air gestures. Ubicomp 2014.
2. Goel, M., Lee, B., Islam Aumi, M. T., Patel, S., Borriello, G., Hibino, S., & Begole, B. SurfaceLink: using inertial and acoustic sensing to enable multi-device interaction on a surface. SIGCHI 2014.
3. Hinckley, K., Ramos, G., Guimbretiere, F., Baudisch, P., & Smith, M. Stitching: pen gestures that span multiple displays. AVI 2004
4. Luo, J., & Zhan, X. Characterization of Smart Phone Received Signal Strength Indication for WLAN Indoor Positioning Accuracy Improvement. Journal of Networks 9(3) 2014
5. Marquardt, N., Hinckley, K., & Greenberg, S. Cross-device interaction via micro-mobility and f-formations. UIST 2012.
6. Peng, C., Shen, G., & Zhang, Y. BeepBeep: A high-accuracy acoustic-based system for ranging and localization using COTS devices. SenSys 2007.
7. Potortì, F., Corucci, A., Nepa, P., Furfari, F., Barsocchi, P., & Buffi, A. Accuracy limits of in-room localisation using RSSI. APSURSI 2009
8. Qiu, J., Chu, D., Meng, X., Moscibroda, T. On the feasibility of real-time phone-to-phone 3d localization. SenSys 11.
9. Quan, M., Navarro, E., & Peuker, B. (2010). Wi-Fi Localization Using RSSI Fingerprinting.
10. Rädle, R., Jetter, H. C., Marquardt, N., Reiterer, H., & Rogers, Y. (2014, November). HuddleLamp: Spatially-Aware Mobile Displays for Ad-hoc Around-the-Table Collaboration. ITS 2014.
11. Wang, Y., Yang, X., Zhao, Y., Liu, Y., & Cuthbert, L. (2013, January). Bluetooth positioning using RSSI and triangulation methods. CCNC 2013
12. Weiser, M. (1991). The computer for the 21st century. *Scientific American*,*265*(3), 94-104.
13. Zhang, Z., Chu, D., Chen, X., & Moscibroda, T. Swordfight: enabling a new class of phone-to-phone action games on commodity phones. MobiSys 2012