

Point&Connect: Intention-based Device Pairing for Mobile Phone Users

Chunyi Peng¹, Guobin Shen¹, Yongguang Zhang¹, Songwu Lu²

¹Microsoft Research Asia, Beijing, 100190, China

²Univ. of California, Los Angeles, CA 90095, USA

{chunyip,jackysh,ygz}@microsoft.com, slu@cs.ucla.edu

ABSTRACT

Point&Connect (P&C) offers an intuitive and resilient device pairing solution on standard mobile phones. Its operation follows the simple sequence of *point-and-connect*: when a user plans to pair her mobile phone with another device nearby, she makes a simple hand gesture that points her phone towards the intended target. The system will capture the user's gesture, understand the target selection intention, and complete the device pairing. P&C is intention-based, intuitive, and reduces user efforts in device pairing. The main technical challenge is to come up with a simple system technique to effectively capture and understand the intention of the user, and pick the right device among many others nearby. It should further work on any mobile phones or small devices without relying on infrastructure or special hardware. P&C meets this challenge with a novel collaborative scheme to measure maximum distance change based on acoustic signals. Using only a speaker and a microphone, P&C can be implemented solely in user-level software and work on COTS phones. P&C adds additional mechanisms to improve resiliency against imperfect user actions, acoustic disturbance, and even certain malicious attacks. We have implemented P&C in Windows Mobile phones and conducted extensive experimental evaluation, and showed that it is a cool and effective way to perform device pairing.

Categories and Subject Descriptors

H.5.2 [User Interfaces]: Interaction styles; B.4.3 [Input/Output And Data Communication]: Interconnections; C.2.1 [Network Architecture and Design]: Wireless communication; C.3 [Special-purpose and Application-based Systems]: Signal processing systems

General Terms

Design, Human Factors, Performance

Keywords

Device pairing, Identification, Gesture-based intention, Acoustic position sensing, Point, Mobile system

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiSys'09, June 22–25, 2009, Kraków, Poland.

Copyright 2009 ACM 978-1-60558-566-6/09/06 ...\$5.00.

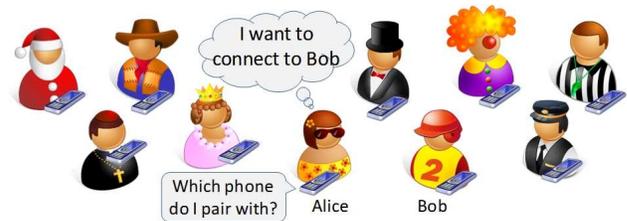


Figure 1: Motivating scenario

1. INTRODUCTION

In recent years, mobile phones have become increasingly popular. This has led to many new applications such as file swapping, music sharing, and collaborative gaming, where nearby users engage in spontaneous wireless data communications through their mobile phone Bluetooth or WiFi interfaces. A prerequisite for such in-situ device-to-device connectivity is *device pairing* – the first-time introduction and association between two devices, often without prior context. Such a connection must be set up before the two phones can engage in an interaction like the above applications.

In this research, we focus on device pairing in a multi-party setting – when there are many mobile devices within the communication range, a user who initiates device pairing must first identify the intended target and convey her selection to the initiating phone in her hand. Figure 1 illustrates a motivating example. Alice has made a new friend Bob and wants to send some pictures from her phone to his. To pair her phone with his, Alice needs a mechanism to let her phone know which nearby phone is the intended target. This involves bridging a “perception gap” in device pairing, where a user knows clearly in her mind what the intended target is (e.g., the mobile phone in Bob’s hand), but she must translate it into a piece of identification information understood at the device level. Currently there are many ways to achieve this. For example, Bluetooth adopts a “scan-and-select” model, where the initiating device scans the wireless channel and lists a set of nearby devices for the user to pick her selection. Other research proposals rely on both parties sharing some private information [10, 18, 25] or taking some synchronized actions together [11–13, 16]. All of them require various degrees of user involvement and system effort.

In this paper, we seek to further minimize this perception gap with a new intention-based device pairing paradigm. It is based on the ability for a user to express, and the system to capture, an intention of device selection via a simple pointing action. The proposed solution, *Point & Connect* (or P&C for short), works as follows. When a user wants to connect her phone to another device, she can express her intention by simply *pointing* her mobile phone towards the intended target device (see the example in Figure 2). The mo-

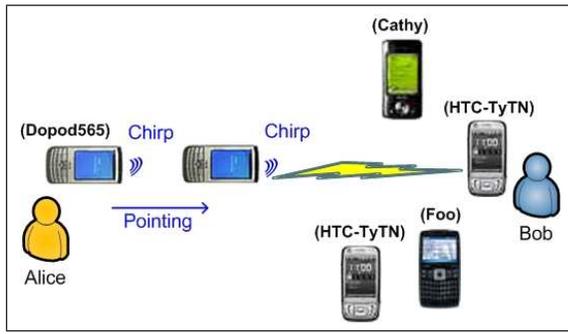


Figure 2: P&C example: Alice pairs her mobile phone with Bob's

mobile phone software can capture this intention and select the right device to complete the pairing operation.

The fundamental challenge here is how to effectively capture and understand the intention of such user pointing action. That is, among the many nearby devices, the system must be able to identify the one that the user's mobile phone is pointing to. It is a non-trivial problem to establish the positioning relationship among devices in the 3-D physical world. Approaches such as visual recognition [19] or applying motion tracking [14] either have too much overhead or require an infrastructure which is not available for in-situ use.

We have met this challenge with a novel acoustic-based distance change detection technique. When a user moves her phone towards the target device, the system requires each "candidate" nearby (at least the target device) to observe and report the relative distance change. To assist the measurement, the selecting phone will emit two sequential "chirp" sound signals during the pointing action and each candidate will independently compute the elapsed time of arrival (ETOA) of the two chirps heard. P&C subsequently exploits the fact that the candidate along the pointing direction should report the largest relative distance change. In addition, P&C devises several techniques to enhance its robustness against various dimensions of factors, including user operational uncertainty and ambient noise, and to defend against a few common malicious attacks.

Of equal importance, P&C is easy to implement on mobile phones. The above novel mechanism operates on a minimum set of hardware capability – the built-in speaker and microphone. This commodity-based solution will obviously have wider applications and cost less, because speaker-and-mic is indeed quite often a common denominator of many mobile devices, including mobile phones, PDAs, media players, etc. In addition, P&C can be implemented solely in software and in user-space, making it easy to adopt and deploy.

The rest of this paper is organized as follows. Section 2 identifies the challenges and reviews the solution space. Section 3 provides the P&C system overview. Sections 4 and 5 present the intention-based design and sensing techniques, respectively. Section 6 describes techniques to enhance resilience against uncertainty and malicious attacks. Section 7 presents the prototype implementation and Section 8 evaluates the performance. Sections 9 and 10 discuss additional issues and compare with the related work, respectively. Section 11 concludes the paper.

2. DEVICE PAIRING MODELS

2.1 Design Challenges

Device pairing typically involves two steps, target identification and secure communication channel setup. The second step is well studied and there are many successful approaches, but most device

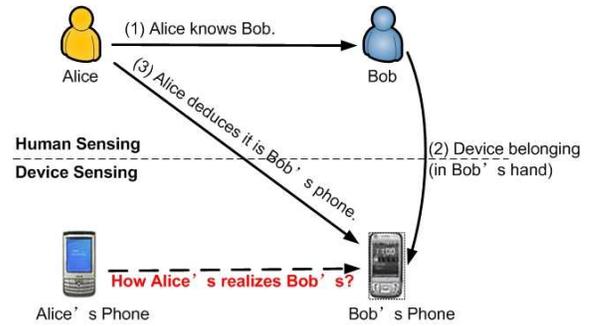


Figure 3: Perception gap between user and device space in device pairing

pairing solutions ignore the target identification issue and leave it for users to handle manually.

The main challenge in target identification is how to bridge the perception gap between how a human identifies a nearby device in mind and how devices must identify themselves among one another, as illustrated in Figure 3. On one side, it is natural for a person to identify a surrounding device by its relative position. Using the same motivating example, Alice knows which device she wants to talk to and where it is – the mobile phone in Bob's hand. On the other side, device pairing requires device-level identification information, such as device name, MAC address, or some other alias, which may easily confuse the user since many of them do not use an intuitive, easy-to-identify naming convention. To bridge this gap under current device pairing models, one must translate the entity in her mind to the device-level identification. In the Bluetooth example, Alice would have to ask Bob for his phone's device name and look it up in the list returned by device discovery. Even if Bob takes effort and discovers his device name (e.g., 'HTC-TyTN' in Figure 2), Alice may still face the confusion: which name corresponds to Bob's phone as there happen to be two devices, both taking the default device name 'HTC-TyTN'.

Ideally, it is desirable to have a device pairing paradigm without such a perception gap. That is, a user needs only to convey her selection target in mind, perhaps with some form of natural action, and the system can capture the user intention and automatically identify the right pairing target. The technical challenge for this vision is to find an intuitive and yet effective way to express an intention of selecting a device, and at the same time to ensure an efficient system can be built to capture and understand such an intention. The engineering challenge is to build such a system that works under current hardware and software constraints, preferably on COTS mobile phones and with only user-level software, and yet be resilient against imperfect user actions, ambient disturbances, or even some malicious attacks.

2.2 Other Approaches

Many existing device pairing solutions adopt a model similar to Bluetooth's "scan-and-select". To bridge the perception gap, a user will need to manually map the intended target into some form of IDs, device names, or network addresses.

If two mobile phones share some private knowledge, this can act as the basis for device selection. Much prior work such as Seeing-Is-Believing [18], blinking-LED [25], Loud-and-Clear [10], and others [28,31], etc., can be used here to assist device selection, even though they are originally developed for device authentication.

Two users interested in device pairing can use explicit synchronized actions as the basis for device selection and subsequent pair-

ing [11–13, 16, 22]. For example, two users can make synchronous actions of button press and release on both phones [22], or place their phones together and move in same trajectory [12]. These approaches usually require both users to agree upon a common action sequence, and some further require mobile phones to be equipped with special sensors such as motion sensors used in [12].

P&C belongs to yet another category in the device pairing solution space. The selection intention is expressed through hand gestures which are simple, intuitive, and very natural.

3. P&C DESIGN OVERVIEW

P&C offers a pure software-based solution to device pairing, and works with a standard phone hardware setting of a microphone, a speaker, and a wireless interface such as Bluetooth and/or WiFi. It does not rely on any infrastructure support or any operating system (OS) modifications. The wireless connection, thus established between the device pair, works in an ad-hoc, peer-to-peer fashion.

3.1 Design Guidelines

P&C should explore an intention-based design paradigm taking the action sequence of point-and-connect. A user will express her intention via actions, and user action and device software will work together to overcome the hardware limitation of a COTS phone. Specifically, the function of target device identification will use the simple human gesture of “pointing”.

P&C should further provide a verification function that exploits human perception to validate the selection result. For example, humans can easily see a blinking screen or hear a chosen sound such as “horn” on the selected device. These natural human inputs can greatly simplify the design complexity on the device.

P&C should also seek to balance operation simplicity and system resilience. It is resilient under a variety of imperfect operating environments. It functions well in the presence of environmental noise, imperfect human gesture, and multipath fading over the acoustic frequency band. It can also defend against a few malicious attacks. However, we seek to reduce the design complexity whenever possible. To this end, we should only use a few fairly simple techniques to enhance resilience, including leveraging human perception feedback, updating blacklists for attackers, and adaptive backoff. For example, in the scenario of Figure 2, if Alice erroneously points to a wrong target, say, Cathy’s phone, then Cathy’s phone will blink its screen or emit a chosen, alarm sound once selected. Alice can simply perform the pointing gesture again. Nevertheless, we do not make it a goal to offer a maximum degree of security against the largest set of possible attacks, because doing so will increase design complexity and reduce its usability and simplicity.

3.2 Application Scenarios

The application scenario for P&C is fairly generic. It is applicable to simple file sharing, picture swapping, and instant messaging. It is also applicable to social networking applications via phones. The typical setting for a P&C working scenario is that the pair of devices, which intend to inter-connect, are located within a short distance. Therefore, line-of-sight communication over the acoustic channel is possible. The devices do not need multi-hop wireless communication for data exchange. Besides the scenario illustrated in Figure 2, P&C also works when Alice selects a device out of many devices that are *statically* placed (i.e., not actively carried by users). The setting is a showcase of scenarios at home, conference room, classroom, etc., where the user intends to pair the device in her hand with another device already in a permanent location.

P&C operates when the relative positioning of devices does not change much. It still works when the devices are moving together,

such as all of them are in a moving bus or train. It does not focus much on the device mobility case, where a few devices are highly mobile relative to other devices. P&C targets a single pair of devices that initiate interconnection at a given time. Its base design does not support multi-pair, simultaneous pairing in a given application scenario. In such a case, other pairs of devices can simply wait until the current pair finishes their connection process which, in any case, will typically last only a few seconds.

4. INTENTION-BASED DEVICE PAIRING

Our intention-based device-pairing paradigm employs two main mechanisms. First, the pointing gesture is used as a way to express user intention in target selection. This nameless pairing bridges the gap between user perception and device naming. Second, the detection of the intended target is achieved through a collaborative distance change measurement method. In this section, we discuss both mechanisms in detail.

4.1 Intentional Selection via Pointing Gesture

The first issue is how to let a user express, and make the system capture, her intention to select a given target device. The solution has to be simple and intuitive, in order to work well with an ordinary user.

Our solution is to let the user make a simple gesture of “pointing” the device to the target phone. The user holds the cell phone in hand, and performs a “pointing” gesture that moves her arm, as well as her phone in hand, straight towards the target device. The rationale for this design comes from natural human behavior in daily life – it is very natural for users to identify and select objects via pointing-style hand motions. This philosophy is also supported by a recent user-behavior study on mobile phone interaction in literature [23]. The study shows that the most common user-object interactions are the actions of *touching*, *pointing*, and *scanning* when performing user-mediated object selection and indirect remote controls. When the user and the object have a line of sight but still are separated by a distance, pointing is the preferred choice.

The above solution addresses the issue of the perception gap between the user and the device. The users do not need to know any device naming. The target device is implicitly identified with its current physical location, and the selection is expressed through the pointing action. In a sense, P&C offers a nameless device-pairing scheme since the semantic name of the device, e.g., MAC address, IP address, or alias, which are in general difficult for non-technical people to understand, is not needed in the selection process.

We believe our solution further reduces user effort needed in device pairing, compared with various other previous approaches described earlier. For example, we do not require precise synchronized actions from both users. Further, our solution does not require motion sensors or other special hardware.

4.2 Detection via Maximum Distance Change

To capture the user intention of a pointing action in a simple and effective manner, P&C devises a simple detection technique based on *relative* positioning. The main idea is based on a geometric, triangle inequality regarding the relative distance change when the pointing action is taken. The example in Figure 4 best illustrates how the detection works. In the example, Alice selects Bob’s mobile phone at position B , and Cathy’s phone is also located at a neighborhood position C . To this end, Alice moves her phone straightly towards Bob’s from position A to position A' . After this pointing gesture is completed, Bob’s and Cathy’s phones see distance changes equal to $d_{AB} - d_{A'B} = d_{AA'}$ and $d_{AC} - d_{A'C}$, respectively. Regarding these two distance changes, the triangle in-

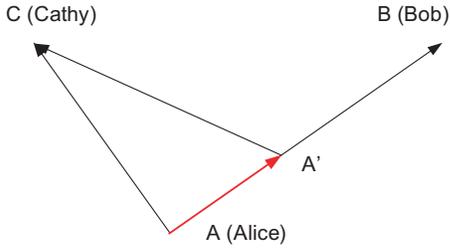


Figure 4: Principle of the P&C detection mechanism. As Alice points her phone towards Bob’s, Bob sees a larger distance reduction than Cathy, as ensured by the triangular inequality: $d_{AB} - d_{A'B} = d_{AA'} > d_{AC} - d_{A'C}$.

equality states that, $d_{AA'} > d_{AC} - d_{A'C}$. That is, Bob’s phone (the intended target device being pointed to) will see the maximum distance reduction before and after the pointing action.

In summary, P&C correctly selects the pairing target by measuring and comparing the distance change at each candidate device before and after the pointing action. Among all the candidates, the device being pointed to will observe the maximum distance reduction before and after the gesture, as ensured by the well-known triangular inequality. Therefore, P&C infers the device orientation associated with the pointing action through detecting the relative distance change at each candidate device. The direct relative positioning scheme further removes the dependency on high accuracy localization infrastructure.

The P&C detection mechanism also exhibits benefits compared with other design alternatives. The popular approach to detecting the pointing action is to use visual markers such as QR code [2], Semacode [3], or image processing [27] with the assistance of a camera. This approach needs cameras and requires the detected objects be tagged with visual markers in advance. It is computationally intensive, thus not suitable for mobile phones and other COTS embedded devices. Moreover, the practical constraints in placing cameras and markers also limit the operating scope. Techniques using laser or infrared pointers can detect a direct pointing action with a lesser degree of difficulty. However, such techniques require that objects be placed within a short range and need special hardware such as a light sensor or an IrDA, which is not presently available on most mobile phones.

5. ACOUSTIC SENSING IN P&C

To make the intention-based scheme work, we need to devise a sensing solution that accurately detects the distance change at the target, which is surrounded by many devices. The solution has to work with the standard hardware setting of a COTS mobile phone. To this end, we need to address three technical issues in P&C sensing: (1) How to achieve high accuracy relative distance change measurement without any pre-deployed localization/ranging infrastructure? (2) How to perform scalable measurement such that the measurement overhead remains constant as the number of candidate devices increases? (3) How to find/notify the intended target, out of many candidate devices, with as little communication overhead as possible, without any *a priori* one-one communications established with the target?

5.1 Detecting Distance Changes

To detect the relative distance changes, several alternative solutions are available but all have downsides. A simple way is to detect the relative difference of received signal strength, say, RSSI values,

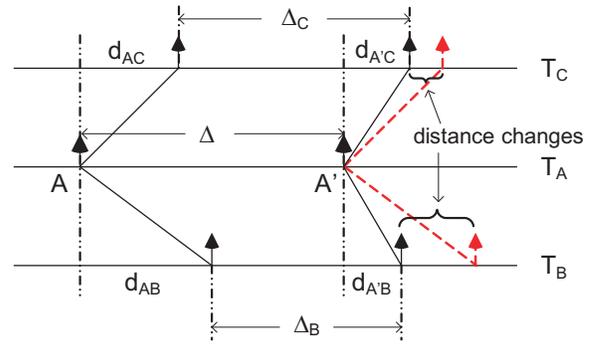


Figure 5: Illustration of the translation from maximum distance reduction to minimum interval between the two sound signals.

at different receiving devices. However, two factors severely degrade the accuracy of this RSSI-based scheme. First, RSSI values fluctuate a lot (e.g., 3 – 8 dB according to previous measurements) even at a fixed receiver location. Second, the distance difference from the sender to each receiving device also makes the distance base for comparisons different. Another way to obtain such distance changes is to derive them from the sequential, pairwise ranging measurements taken before and after the pointing action. However, since pairwise ranging is performed between the phone and each candidate target, it does not scale: To select one out of N devices, $2N$ ranging measurements should be performed (to obtain the distance changes) and the overall process would be extremely time consuming and thus more vulnerable to ranging errors. The accuracy is impaired as the error bound is doubled because of the distance subtraction operation.

We propose a novel solution that eliminates the scaling issue incurred by the sequential scheme. We let the selecting device emit two sound signals before and after the pointing action, and ask all the candidates to record and detect the two signals and report the interval in between, then we can translate maximum distance reduction to the minimum measured interval and select the winner accordingly.

Figure 5 illustrates such a translation. Because all the devices are located in a close proximity, the propagation speed of sound signal to all candidates will be the same. The distance is therefore strictly proportional to the time duration the sound signal propagates, as depicted by the slanted lines in the figure. If the selecting device does not move, then the second sound signal will take the same time as the first signal to reach the candidate devices as indicated by the two dotted slanted lines. In this case, the measured intervals at candidate devices, Δ_C and Δ_B , will be equal to Δ – the interval between the two sound signals at the selecting device. As the selecting device moves closer to the candidates after the pointing action, it takes less time for the second sound signal to reach them. As seen in Figure 5, the distance reduction at B can be expressed by

$$\Delta_{dB} = d_{AB} - d_{A'B} = v \cdot (\Delta - \Delta_B), \quad (1)$$

where v is the sound speed and is also roughly the same in close proximity. The distance change at C , smaller than that at B , can be measured similarly. Therefore, the maximum distance reduction can be translated to the minimum detected interval.

Therefore, the problem (1) of high accurate relative distance measurement is converted to obtaining the accurate, elapsed time of arrivals (ETOA) of two sound signals in the presence of an inac-

curate system clock. The traditional way of taking a timestamp of their respective local clocks at the moment the signal is received to calculate the duration may meet with intrinsic receiving uncertainties due to the lack of real-time control, software delay, interrupt handling delay, and system loads in a real system. Therefore, we use the sampling counting scheme proposed in [20] that addresses the problem. The solution leverages the fact that the independent A/D convertor works at a fixed and stable sampling frequency, thus generating high-accuracy time information without unknown delay factors from the phone’s operating system. In P&C, the target device can be determined through comparing the sample numbers between two sound arrivals among multiple devices. The device with the minimal number of samples between two received signals is the one along (or closest to) the movement direction.

Our solution simultaneously estimates the pointing phone’s distance changes to all candidates. No matter how many candidate targets are around, only two signals are required. These two chirps will be used by multiple candidates to estimate the distance changes in parallel. The overall signaling overhead is independent of the number of candidate devices, and our sensing scheme yields an accurate device identification, scalable to the number of candidates.

5.2 Exchanging ETOA Measurements

The above sensing scheme enables a scalable solution to detecting the relative distance change at each candidate device when the pointing action is taken. The next issue is to find and notify the device that detects the maximum distance change, i.e., the intended target. The challenge is to exchange ETOA measurements between the selecting device and the target device, as well as other nearby devices. Our solution is to establish a broadcast channel by exploiting the radio interface on COTS phones, such as Bluetooth or WiFi. We also make an effort to avoid potential contentions (and conflicts) over the broadcast channel.

Two application scenarios occur in practice. In the first case, all devices are already in a wireless network (e.g., WiFi), but they do not know which device is the target one. This scenario appears in settings of office, home, classroom, and cafe with wireless access. In the second case, there is no established radio channel before P&C. It may occur in a spontaneous pairing scenario, where WiFi network setup or Bluetooth pairing is needed. We address both cases.

In the first case, we simply let all candidates listen to the current, pre-determined wireless channel and report their ETOA measurements over it. We use a backoff-based mechanism to detect the target device, while avoiding potential contentions over the channel. Each device sets a backoff timer, in proportion to its ETOA interval measurement. The backoff timer is implemented at the link buffer, thus different from the low-level MAC backoff mechanism. The timer granularity is also coarser, say, tens or hundreds of milliseconds for each backoff unit in our implementation. The backoff scheme effectively provides a fully distributed mechanism to prioritize the reporting process based on the increasing order of ETOA measurements. We also include a fixed offset (e.g., 300ms in our implementation) to the backoff timer to absorb the time differences for ETOA detection on different phones. Each device monitors the broadcast channel and ensures a clear channel before sending out any message. Therefore, using the backoff timer, the selecting device can quickly find the target phone by receiving messages with small ETOA values first. Ideally, the target device, which has minimum ETOA, will send its message (ETOA measurement and its MAC address) first, and others will defer when hearing such an early message. Upon receiving and identifying the device that reported minimum ETOA, it subsequently sends an “acknowledg-

ment” message to notify the device. A one-to-one data communication channel can then be set up between the pairing devices.

In the second case, we use the *Namecast* method to bootstrap and establish a broadcast radio channel. Namecast is a hack of the typically neglected broadcast channel, which is used to announce the device name (e.g., the Bluetooth device name or the SSID field of WiFi beacons) and enable device discovery. It works as follows. Each candidate temporarily overloads its device name to embed the desired information (e.g., the ETOA measurement and its MAC address); the selecting device can effectively retrieve such information by scanning Bluetooth device names or WiFi beacons. A similar concept has been explored in Client Conduit [7]. Namecast works for both Bluetooth and WiFi. The downside of Namecast is, however, the relatively long scanning time, which incurs one-time cost during the channel bootstrapping process.

Note that our above solution also scales to the number of candidate devices. The selection process does not incur additional ETOA message overhead, in that the number of transmitted ETOA messages does not grow with the device population.

6. ENHANCING RESILIENCE

In this section, we first identify a few factors that may affect the resilience of P&C and then discuss the proposed solution techniques. The goal is to ensure that P&C is robust against various uncertainty factors. While P&C design is targeted at applications for which security is of not critical concern, we still present a few techniques that helps to improve the security against a few common, malicious attacks.

There are various dimensions of uncertainty in the pointing action, device, ambient disturbances, as well as malicious attacks. These factors can be roughly classified into the following three categories:

- Imperfect user action: various gesture uncertainties or operation errors may exist.
- Acoustic sensing disturbances: various error sources exist, including multipath propagation and fading, ambient noise, etc.
- Malicious attacks: practical attacks may occur, such as denial-of-service (DoS) attacks, and man-in-the-middle (MiTM) attacks.

We now describe how P&C handles the above factors in detail.

6.1 Handling Imperfect User Action

The first type of error is related to human gesture uncertainty when using P&C. There are three common errors/mistakes a user makes during the pointing action: (1) pointing to the wrong target: Instead of pointing to the target device, the user mistakenly points to other objects or devices located in the neighborhood. (2) non-direct pointing action: The pointing trajectory by the user is not a straight line towards the target. Rather, it is a curve or other arbitrary shape. This reduces the effective distance change to the target and may increase the distance change detected at other nearby devices; (3) short pointing: The pointing gesture follows a straight line to the target, but travels too short a distance. This can also compromise the detection accuracy in P&C.

Our solution to the above human gesture uncertainty still takes the user-centric approach. We rely on user perception feedback to resolve the issues. Once selected, the device blinks its screen for a short period of time or an alarm sounds. By observing which device’s screen blinks or sounds the alert right after the pointing action, the user can verify whether she has selected the correct target.

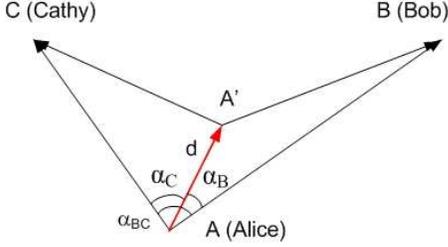


Figure 6: Robustness against pointing orientation error.

If not, the user simply repeats the pointing action to her intended target until she sees the right device blinking. Therefore, the user has an independent source to reliably verify whether the selection is correct or not and correct her mistakes.

Robustness against pointing orientation error We now show that P&C itself is fairly robust against the pointing direction errors. Therefore, the above technique is not invoked in normal cases.

We use the example scenario of Figure 6 to show that P&C is still able to identify the correct target B even when the pointing error α_B is nonzero but not too large, say, less than 45° . By applying Cosine Theorem, the distance reduction at B is:

$$\Delta_{d_B} = d_{AB} - d_{A'B} = b - \sqrt{b^2 + d^2 - 2bd \cdot \cos(\alpha_B)}, \quad (2)$$

where b and d stand for d_{AB} and $d_{AA'}$. The above function is monotonically decreasing with respect to $\alpha_B \in [0, \pi]$, and reaches the maximum peak at $\alpha_B = 0$, independent of b . In particular, when $b \gg d$ and α_B is sufficiently small (the typical case in reality), the reduced moving distance Δ_{d_B} can be approximated by

$$\Delta_{d_B} \approx d \cdot \cos(\alpha_B). \quad (3)$$

Hence, the distance reduction approximates the projection of the “pointing” onto the original direction towards the target device.

We now assess the impact of an imperfect pointing angle α_B (i.e., $\alpha_B \neq 0$) upon the distance reduction in the example scenario, where B and C are fixed at $d_{AB} = d_{AC} = 100\text{cm}$ and $\alpha_{BC} = 90^\circ$ ¹. Figure 7 plots the difference of distance reduction $\Delta_{d_B} - \Delta_{d_C}$ under different pointing movements of 20cm, 30cm and 40cm. The figure shows that, in realistic cases where the pointing error (α_B is small but nonzero, say, less than half of α_{BC}) exists, the one with smallest α (i.e., the one closest to “pointing” direction) can still be correctly identified by comparing the distance reduction.

6.2 Handling Acoustic Sensing Disturbances

The ambient noise and multipath effect may distort or attenuate the acoustic chirp signals, thus significantly reducing ETOA detection accuracy if not handled well. P&C adopts two similar techniques in [20] to handle both background noise and multipath effects, and proposes an efficient joint detection.

To suppress the effect of ambient noise, the overall scheme takes a correlation-based detection approach. Each device records the received signal, which is correlated with the reference chirp signal stored at each device. The maximum peak is located if the cross-correlation of the signal and the reference is beyond a threshold ratio compared with the cross-correlation of the background noise and the reference. To this end, we compute the L_2 -norm of the cross-correlation value within a small window of samples around

¹More cases of different d and α_{BC} , with different combinations of d_{AB} and d_{AC} , yield the same conclusion.

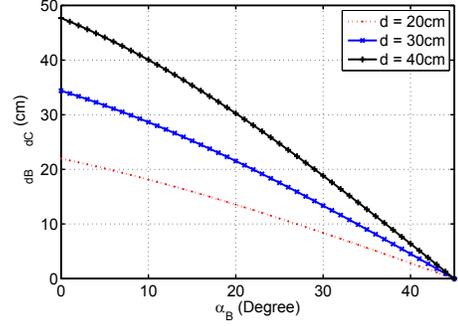


Figure 7: Differential distance change in response to α_B

the peak; this reflects the energy level of the received signal. We also calculate the L_2 -norm of correlation values over a time window right before the peak, which reflects the noise level. When the energy level of the signal is twice or larger than the noise, we infer that the received chirp signal is located and ETOA is then calculated.

To handle the multipath effect caused by reflection from a secondary path, we locate the first correlation peak rather than the maximal one. This is because the maximum peak may appear along the secondary path, which lags behind the signal on the primary path. In P&C, we locate the earliest, sharp peak in the time window of interest. The sharpness captures the level of the peak with respect to its surrounding side-lobes, and is computed by the ratio of the peak correlation and the average cross-correlation values in its vicinity.

The third technique is a joint ETOA detection scheme that can further improve the detection accuracy and also reduce the computation workload. It is based on the observation that P&C uses a pair of chirp signals in its sensing and these signal arrivals follow certain time patterns. We can then detect both arrivals in a joint fashion rather than independently. The elapsed time between two chirps must fall into the range $[\Delta - t_{max}, \Delta + t_{max}]$, where t_{max} denotes the spread time for the maximal pointing distance and Δ is the original interval between the two chirps at the selecting device. Therefore, when the first chirp signal is detected, we only need to further detect the second one in the small window ($2t_{max}$ in width) that is Δ away from current spot. We will conclude a successful detection only when the second signal is also detected in that window. Doing this way also reduces computation since most of samples between two chirps need not be checked and are safely skipped. The computation savings (for the second chirp signal) can be significant in practice, up to $1 - 2t_{max}/\Delta$. This evaluates to more than 99% savings for practical settings with 44.1kHz sampling rate, one meter pointing distance, and one second interval between two chirps.

6.3 Defending Against Malicious Attacks

DoS attacks include generating similar chirp signals, and forging a small ETOA value. The attacker can also perform a MiTM attack over the wireless channels.

Forged Chirp attack This is one type of DoS attack. The attacker plays one or more similar chirp sounds once he hears the first chirp signal emitted from the selecting device. Using a forged chirp signal may collide with the second chirp signal emitted from the selecting device, or confuse the receiver(s) so that for an extended period of time, the attacker effectively disrupts the reception of the second chirp signal from the selecting device. We take

a reactive solution. We conclude a failure when more than two chirps are heard during a reasonable time frame of a pointing action (e.g., two seconds in our implementation). In essence, the solution relies on the user to partially defend against such an attack. However, there is no perfect solution to this type of physical-layer DoS attack. Nonetheless, the aforementioned joint ETOA detection mechanism can effectively mitigate such attacks since the forged chirps are not effective unless they happen to be emitted within the $[\Delta - t_{max}, \Delta + t_{max}]$ range.

Fake ETOA attack This is another type of denial of service (DoS) attacks against P&C. In the attack, the malicious attacker fabricates a smaller ETOA value than it should be, in order to cheat the user into never selecting the right target but instead choosing the attacker’s device. We use three means to combat such attack. First, P&C filters out any invalid ETOA value that is beyond the typical operation range. If the attacker without hearing two chirp signals sets an overly small value, it cannot pass the ETOA filter. This provides the first line of defense against a forged ETOA attack. Second, the selected target is asked to blink his screen or alarm an alert once selected. Then the selection failure of choosing a wrong device (the attacker’s phone in this case) is quickly detected by the user via visual observation. Once detected, the user will re-take the P&C operations by pointing to the target device again. Third, P&C records the attacker device’s ID and places it onto the *blacklist*. A device on the blacklist will be forbidden to be selected for a period of time, say, 5 minutes. Note that the blacklist method also prunes away potential wrong candidates, which may result from users’ imperfect actions.

MiTM attack The attacker can impersonate each endpoint to the other, making the victims believe that they are talking directly to each other when in fact they are talking with the attacker independently. A common solution to prevent MiTM attacks is endpoint authentication. P&C can leverage existing device authentication solutions [10, 28]: it uses the acoustic channel to exchange both data and verification information among devices and involve users to verify the same audio source from the right device pair.

7. IMPLEMENTATION

P&C is implemented as a user-level software solution. In this section, we describe the work flow of P&C and implementation details.

7.1 P&C Work Flow

The overall procedure of Point&Connect is as follows:

Init: All participating devices pre-set up the system and start the P&C service. Detailed tasks include (1) preparing the chirp signal; (2) turning on the radio interface, and tuning to a common channel (WiFi) or discoverable mode (Bluetooth); (3) turning on the microphone and starting recording.

Point: The selecting user makes a “pointing” gesture. It is triggered by a simple activation button on the phone. When the button is pressed, the selecting device emits a chirp signal. The user holds the button and moves toward the target device until the second chirp signal is emitted, at which point the button is released and the pointing gesture is complete—release the button then and complete the pointing gesture.

Sense: This step is executed at all other nearby devices, including the target one during a “pointing” action. It records the sound signals over the acoustic channel, and thus captures

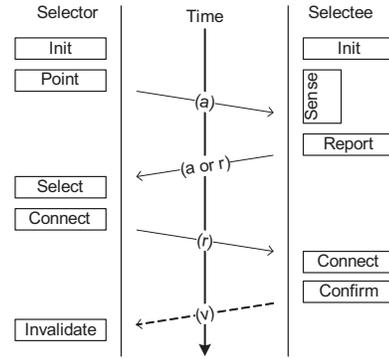


Figure 8: Block diagram of P&C, with temporal activity ordering. Characters ‘a’ or ‘r’ striking the arrows indicates the information is conveyed using acoustic or radio interface; and ‘v’ on dashed line implies through visual or vibration feedback.

both chirp signals. Once each receiving device detects the two chirp signals, it stops recording. Then it calculates the elapsed time (measured in the sample counts) between the two chirp signals.

Report: Each device reports its measured ETOA value and also its MAC address, by changing their device name (for Bluetooth) or over the common broadcast WiFi or acoustic channel. For the latter case, the backoff-based mechanism is also used.

Select: The selecting device scans surrounding devices’ names (for Bluetooth) or listens to WiFi channel to retrieve reported ETOA values and their associated MAC addresses. It then identifies the intended target device.

Connect: The selecting device connects to the identified target using the MAC address (Bluetooth) or joins the WiFi network with the target’s SSID (WiFi), and the target device in return accepts the connection request and sets up a private radio communication channel.

Confirm: The selected device flashes its screen or vibrates itself to confirm the selection. This step provides another independent source of feedback to users. It is informative and provides better usability.

Invalidate: The user may invalidate the connection when a wrong device is selected. When it is invoked, the module tears down the radio communication channel with the previous device that was set up during the last connect() operation.

7.2 Implementation Details

In our implementation, P&C has eight function modules, each corresponding to one step in the work flow. The eight modules work in concert in the current prototype. Figure 8 shows a block diagram of how they work during the P&C operations. The three modules *sense*, *report* and *confirm* are major functionalities for the selectee. The other three modules *point*, and *select*, *invalidate* are for the selecting device. The *Init* and *connect* modules are used by both.

The above modules have been implemented on smartphones and PocketPC phones running Windows Mobile 5.0. We use multimedia services (*waveXxx* series APIs) embedded in Windows Mobile, to control microphones and speakers on the phones; and use the NDIS User Mode I/O interface (NDISUIO) and the Windows Sockets functions (*WSALookupServiceXxx*) and structures for

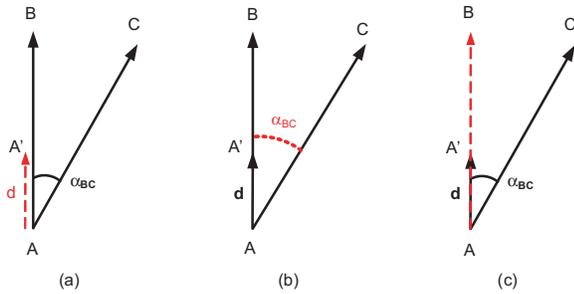


Figure 9: Experiment settings. We vary one factor at a time and keep the remaining factors unchanged in each experiment. Broken and thick solid lines indicate the varying and fixed factors in each test case.

wireless communication. The software runs as a user-mode dynamic linkable library that other applications can load and use as a service. We have deployed the applications on several smartphone platforms including Dopod 838 (i.e., HTC Wizard), HTC S620, HTC P3300, HTC Touch, HTC Tornado, HP iPAQ rw6828 and MWG Atom Life etc.

8. EVALUATION

In this section, we use both experiments and numerical analysis to assess the performance of P&C.

We have evaluated our implementation using four phone models, Dopod 838, HTC Tornado, HP iPAQ rw6828 and MWG Atom Life, all running Windows Mobile 5.0 or above. Other phone models produce similar results. In the rest of this section, we assess P&C correctness against different parameters, examine resilience on imperfect pointing, report P&C energy consumption, and describe field trial experiences. For experiments on P&C correctness and tolerance (Section 8.1), we have pre-configured all devices into a WiFi network and exchange ETOA over a wireless channel. For field trials, we experimented on different end-to-end device pairing processes, as detailed in Section 8.4.

8.1 P&C Correctness and Tolerance

We first evaluate the correctness of P&C against factors, which include: displacement of the selecting phone (d) in the pointing action, target separation angle (α_{BC}), and distances between phones. We evaluate one factor at a time while fixing all the other factors. Test cases are depicted in Figure 9. Each experiment is repeated for 20 sequential runs (i.e., randomly selecting a target device out of the two candidates, pointing to it and recording the selection results as indicated by screen blinking) by the same person. Experiments are conducted under different environmental settings, such as quiet indoor (conference room) and noisy outdoor (near subway station entrance) environments. The user always points directly to the target when using P&C, except the evaluation on imperfect user actions.

Phone displacement d This set of experiments examine the impact of displacements d of the selecting phone on the correctness of P&C. We tested the performance under different combinations of α and the distance between phones, as shown in Figure 9-(a). Due to practical constraints (e.g., a person’s arm length), we only evaluate the displacement distance up to 60 cm.

Figure 10 plots the test results. We can see that larger displacement leads to more reliable selection. The results suggest that in practice, we need to move the phone by at least 20cm to obtain

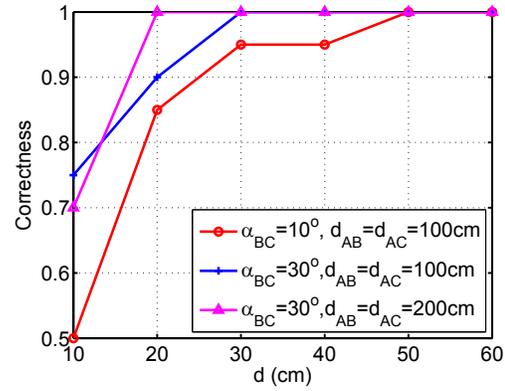


Figure 10: P&C correctness vs the displacement of the selecting phone in the pointing action

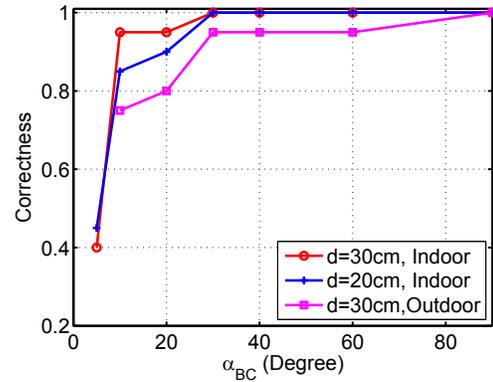


Figure 11: P&C correctness vs the target separation angle. In all cases, we fixed $d_{AB} = d_{AC} = 100\text{cm}$.

accurate selection. Furthermore, when the candidate devices are separated by only a small angle (e.g., 10 degrees), a larger displacement is needed.

For the sake of practical interest, we set the displacement of the selecting phone to 20cm and 30cm in the remaining experiments.

Target separation angle α_{BC} We evaluate the relation between the target separation angle α_{BC} and the selection correctness. In all cases, we have fixed the distance $d_{AB} = d_{AC} = 100\text{cm}$. The results are shown in Figure 11. As expected, the larger the separation between the target device and its closest neighbor, the more accurate the selection is. In most cases, 20cm displacement works fairly well, when the devices are separated by 30 degrees or more.

Note that, however, when α_{BC} is very small, the selecting result suffers and looks more random. This is easy to understand since when α_{BC} is small, it is indeed very challenging even for a person who makes a large-distance pointing gesture. In our setting, $\alpha_{BC} = 5^\circ$ implies the distance between the target and its nearest neighbor is only 9cm. The figure shows that there is also a sudden increase in correctness when α_{BC} increases. This reflects the effect of the ETOA detection accuracy ε_{res} . When α_{BC} is very small, the distance changes at B and C are within ε_{res} and thus cannot be differentiated.

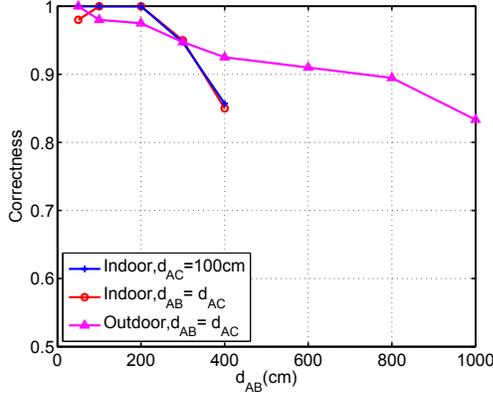


Figure 12: P&C correctness vs d_{AB} . In all cases, we fixed α_{BC} to 30 degrees, and $d = 30cm$.

Distance d_{AB} This set of experiments evaluates the impact of the distance d_{AB} between the selecting phone and the target on the selection correctness. In all experiments, we fixed $\alpha_{BC} = 30^\circ$ and $d = 30cm$. In the first case, we fixed the distance $d_{AC} = 100cm$. In other cases, we let the distance vary while keeping $d_{AC} = d_{AB}$. The results are plotted in Figure 12. We see that the selection is very accurate when the distance is less than three meters. The absolute distances, from the target device or its neighbors to the selecting device, have little impact on the identification.

However, when the distance gets even larger (e.g., four meters in indoor scenarios and eight meters in outdoor settings), the selection correctness degrades. This is not consistent with the theoretical analysis to be presented in Section 8.2. In theory, d_{AB} should have little impact on selection correctness when $d_{AB} \gg d$. We examined the experimental traces, and found that this result is due to the drops in the ETOA detection accuracy, as caused by the multipath effect and signal energy reduction for indoor and outdoor cases, respectively. Similar findings have also been reported in [20].

Tolerance on imperfect pointing angle α_B This set of experiments examines the effect of an imperfect pointing angle α_B (refer to Figure 6) on P&C. In all experiments, we fixed the parameters as $\alpha_{BC} = 60^\circ$, and $d = 30cm$. The results are shown in Figure 13. Indeed, we see that in most cases, we still correctly select the target as long as the pointing direction leans more towards the right target. The result also confirms that the direct pointing action ($\alpha_B = 0$) yields the best results. However, there is a sharp drop when the pointing angle increases from 20° to 30° , which is about half of α_{BC} . This confirms that the pointing error cannot be too big.

8.2 Numerical Assessment on Resolution

In this section, we provide a numerical assessment for the P&C resolution and learn practical ways to improve resolution, which refers to the minimum spatial requirement on neighboring devices so that they can be differentiated from the target device. Note that the P&C base scheme can be summarized as follows:

Target Device	
= the one with minimal α_B ,	\Leftarrow User intention
= the one with maximal Δ_{d_B} ,	\Rightarrow Device iden.
= the one with minimal ETOA,	\Rightarrow Sensing

In the example of Figure 6, the resolution defines the condition

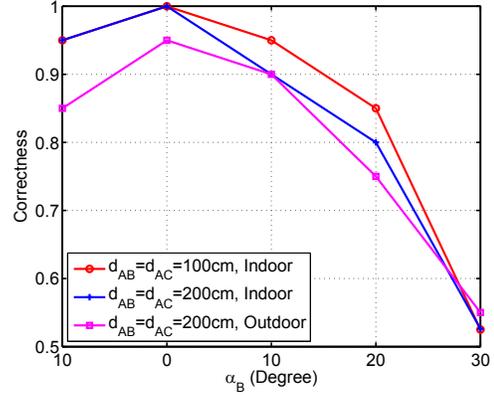


Figure 13: Resilience on pointing error. We fixed the parameters as $d_{AB} = d_{AC} = 100cm$, $\alpha_{BC} = 60^\circ$, $d = 30cm$.

so that B can be differentiated from C . Therefore, the resolution requirement translates to:

$$\Delta_{d_B} - \Delta_{d_C} > d_{res}, \quad (4)$$

where d_{res} denotes the distance reduction measurement resolution. Since we adopted the sampling counting method to measure distance change, d_{res} is thus determined by the ETOA detection accuracy ε_{res} , i.e.,

$$d_{res} = \varepsilon_{res} / f_s * v \quad (5)$$

where f_s is the sampling frequency and v is the sound speed. The minimal d_{res} is achieved at only one sample resolution, and it is roughly $d_{res} = 1/44100 * 34600 \approx 0.7cm$ in a typical setting of $f_s = 44.1KHz$ and $v = 346m/s$. Moreover, BeepBeep [20] has shown that the accuracy d_{res} within 2cm is achievable.

In principle, node C can be differentiated from A when $\Delta_{d_B} - \Delta_{d_C} \geq d_{res}$. While the closed-form derivation can be quite involved, we start with a few realistic cases. For sake of clear expression, let $b = d_{AB}$ and $c = d_{AC}$.

Case 1: Perfect pointing action if $\alpha_B = 0$. In this case, Equation (4) can be simplified as

$$d - \left(c - \sqrt{d^2 + c^2 - 2dc \cdot \cos(\alpha_C)} \right) \geq d_{res}$$

If $c \gg d$, we have

$$\cos(\alpha_C) \leq 1 - \frac{d_{res}}{d} \quad (6)$$

Case 2: The device is far away when $b \gg d$, $c \gg d$. The reduced distances can then be approximated by $\Delta_{d_B} \approx d \cdot \cos(\alpha_B)$, and $\Delta_{d_C} \approx d \cdot \cos(\alpha_C)$. Hence, we have

$$\cos(\alpha_C) \leq \cos(\alpha_B) - \frac{d_{res}}{d} \quad (7)$$

In the setting when perfect pointing to B , $d = 30cm$ and $d_{res} = 0.7cm$ (i.e., 1 sample), only node C with $\alpha_C \subseteq (-10.7^\circ, 10.7^\circ)$ can not be differentiated, which corresponds to about 40cm between B and C that are 2m away from A . In this case, B and C are very close, not a common scenario in practice.

We further plot Figure 14, to show the angle resolution within which close-by nodes cannot be differentiated from the target for

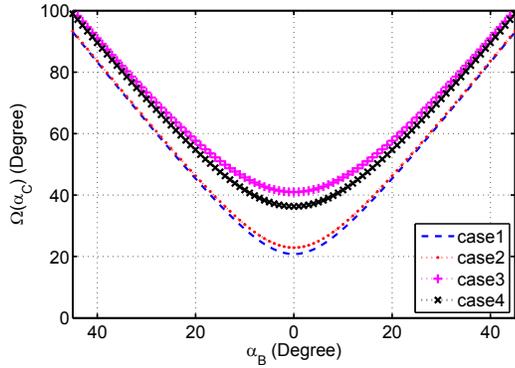


Figure 14: P&C resolution vs pointing angle. Define $\Omega(\alpha_C)$ as $\Omega(\alpha_C) = \max(\alpha_C) - \min(\alpha_C)$, where $\max(\alpha_C)$ and $\min(\alpha_C)$ are the maximal and minimal non-distinguishable angles.

	d_{AB}	d_{AC}	d	d_{res}
Case 1	100cm	100cm	30cm	0.7cm
Case 2	200cm	200cm	30cm	0.7cm
Case 3	100cm	100cm	10cm	0.7cm
Case 4	100cm	100cm	30cm	2.1cm

Table 1: Four example cases for angle resolution

four cases². The four settings are given in Table 1, where we use Case 1 as the main reference for comparisons. From Figure 14, we make several observations on how these factors affect the identification accuracy.

First, the non-distinguishable angle scope increases as $|\alpha_B|$ increases, shown in all cases. It implies that a perfect pointing angle $\alpha_B = 0$ may bring the maximal distance reduction and it is best to distinguish the target device from neighboring nodes.

Second, the absolute distance, from the target node or its neighbor to the selecting device, has little impact on identification. The results at 100cm/200cm are almost the same for Cases (1) and (2). This can also be seen from Equation (7), which intuitively implies that it is the angle, not distance that matters.

Third, the pointing magnitude does matter much when comparing Cases (1) and (3). The larger the movement, the easier to differentiate from neighboring nodes. We can see that $\Omega(\alpha_C)$ decreases as d increases. From the derivation of (7), we have $(\cos(\alpha_B) - \cos(\alpha_C)) = 2 \sin(\frac{\alpha_C + \alpha_B}{2}) \sin(\frac{\alpha_C - \alpha_B}{2}) \leq \frac{d_{res}}{d}$. When $\alpha_C - \alpha_B$ is small, it can be approximated by $k \cdot (\alpha_C - \alpha_B) \leq \frac{1}{d}$. Hence, the resolution scope is inverse proportion to the moving distance. Heuristically, larger d should also increase identification accuracy.

Finally, the sample counting resolution is in proportion to the identification angle resolution, as shown in Cases (1) and (4). The lower accuracy the ETOA detection, the larger the neighbor scope that can not be identified from the target one.

In summary, to improve the identification resolution that separates the target device from other neighbors, we have three options: (1) we can point directly to the target node (i.e., $\alpha_B \rightarrow 0$); (2) we can extend the pointing action (i.e., larger travel distance d when pointing), and (3) we can reduce the sensing granularity (i.e., small ϵ_{res}). These findings are consistent with the experimental results.

²More cases of different combinations of d_{AB} , d_{AC} , d and ϵ_{res} have been tested and yield the same conclusion.

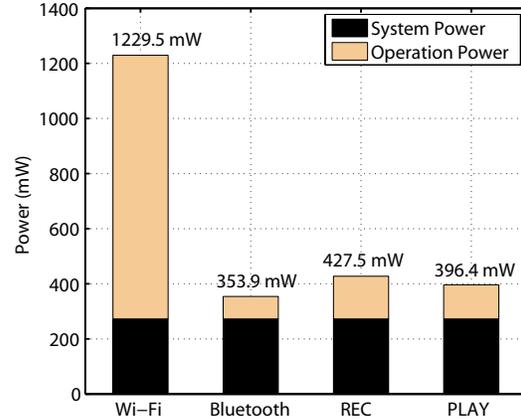


Figure 15: Power consumption of different tasks.

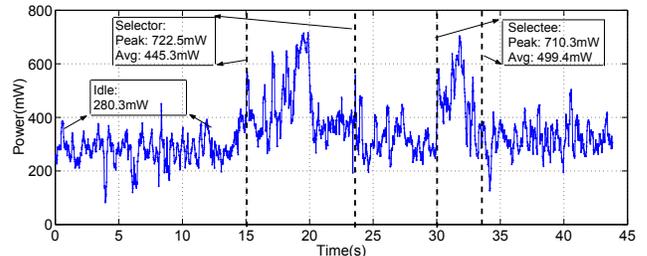


Figure 16: Power consumption of P&C procedure, excluding WiFi energy consumption.

8.3 P&C Energy Consumption

We measure the power consumption of different functional tasks on the HTC Tornado. Figure 15 plots the power consumption of different functional operations. The basic system, including backlight, CPU, and memory, consumes about 275mW. The power consumption for different tasks of WiFi, Bluetooth scanning, microphone recording, and chirp playback is about 957mW, 89mW, 155mW, and 124mW, respectively. We can see that the acoustic sensing operation in P&C, which includes playback of two chirp signals at the selecting device side and recording at the receiving side, needs to use an extra 150/120+mW, comparable to the phone system module. Therefore, a conservative estimate for P&C power consumption is about 430mW on the HTC Tornado phone. However, the WiFi channel is power hungry, whereas the Bluetooth is much more efficient.

Figure 16 shows the power consumption of an entire round of P&C procedures over time, using the WiFi broadcast channel. We did not include the power consumption of WiFi here, to highlight the energy used by main P&C. The first portion ([0, 30] seconds) are for operations at the selecting device, and the second portion ([30, 44] seconds) are for selectees. The procedure at the selecting device includes the initial idle state (when P&C is invoked but no other operation is performed) during [0, 15], and the pointing, selecting ([15, 20]), and idle operations afterwards. The target device starts its recording state for 1-2 seconds, followed by signal detection, ETOA measurements and report etc in [30, 33] seconds. We can see that, compared with the basic system power consumption, a P&C in the idle mode consumes an extra 5mW or so, a P&C as the selector consumes an extra 170mW for about 6-8 seconds, lead-

	Normal	P&C (WiFi)	P&C (NC)
mean	57.2	6.5	27.4
stdev	10.1	1.6	11.4

Table 2: Overall time consumption (in seconds) for Bluetooth pairing with different schemes.

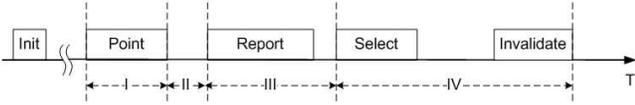


Figure 17: Illustration of time consumption for each step.

ing to 1.2J used energy; each candidate device consumes additional 225mW for about 3 seconds. Therefore, the power consumption by P&C modules is fairly modest, especially when compared with the energy consumption of WiFi. Furthermore, since the overall pairing speed is significantly improved (see next subsection), the actual energy consumption of P&C is actually reduced.

8.4 Field Trial Experience

We also conducted some field trials using our prototype. We invited about 20 college students that have no device pairing experience before. They were divided into 10 pairs and asked to perform normal Bluetooth pairing and P&C Bluetooth pairing. For each pair of students, the two were asked to initiate the device pairing process in turn. The experiments were performed sequentially, and there were no concurrent pairing sessions. Before experiments, we briefed them for 2 minutes, and provided a printout of the detailed steps. We measured the overall time they took to pair up the devices, and recorded the timestamps for each major module. We also asked for their feedback on our prototype. Different phone models are used in experiments, including Dopod838, HTC Touch, HP iPAQ rw6828 and MWG Atom Life models. Since some components (e.g., the pointing gesture) change over different usage patterns, the results also vary.

Table 2 shows the average and standard deviation of the time consumption for a normal Bluetooth pairing process and the proposed P&C with different ETOA exchange mechanisms. We see that the P&C scheme that uses a common WiFi channel for ETOA exchange (which corresponding to the first case scenarios in Section 5.2) is fastest. It takes only 6.5 seconds on average to pair the device and build the logical one-to-one channel. The P&C scheme using namecast takes 27.4 seconds to establish the pairing, which is slower than using the WiFi channel. It is still much faster than the normal Bluetooth pairing process that takes 57.2 seconds on average. Note that our P&C scheme does not improve the device scanning process. The savings mainly come from (1) the simplified user interface as users do not need to follow the lengthy steps in the normal Bluetooth pairing process; and (2) the avoidance of selecting and entering the Passkey at both devices because the user is well aware of the selected device in P&C. We also want to point out that the device scanning process is highly variable, ranging from 10+ seconds to almost 50+ seconds. In consequence, the standard deviation for the normal pairing process and also P&C with namecast are much higher than those without scanning.

We recorded the timestamps for each major operation at the selecting device, as illustrated in Figure 17, to reflect detailed time consumption of user interactions. The time spent on each major operation, averaged over all users, is shown in Table 3, where ETOA is exchanged over the pre-established WiFi channel. We can see

	Mean	Stdev	Min	Max
I	1351.6	958	645	2430
II	366.3	16.1	315.5	393.4
III	2324.7	16.5	2298	2346
IV	2517.6	644	2048.4	3021.1
II+III+IV	5208.5	641.7	4679.4	6102.3

Table 3: Time consumption breakdown of P&C procedures in real experiments by different users, in unit of milliseconds. A common WiFi channel is used for ETOA exchanges.

that, a user typically takes 1.4 seconds to complete the pointing gesture. After 0.4 seconds, the selecting device starts to receive the ETOA report. Note that it actually includes 300ms backoff offset we added in the implementation to accommodate different ETOA detection speeds at different devices. With the backoff mechanism, the selecting device takes another 2.3 seconds to collect more report messages and make the select decision after using the backoff mechanism. Once selected, the phone pair uses about 2.5 seconds to set up the one-to-one communication channel, which includes an average 2 seconds for the user to see the screen blinking and react by pressing a confirmation button in our experiments.

Finally, all the users indicated that our solution is very easy to use and intuitive, with almost no learning curve. This is indeed the key value of P&C, which we think is even more important than the speed improvement.

9. DISCUSSION

When to emit Beeps: In our current design, we have assumed the interval between two chirps used for device selection is fixed and known a priori to all recipients. This requires the user to finish the point action within that interval and may make the user feel a bit rushed. This can be relaxed and enable a more natural user experience. In an alternative implementation, we can let the device emit a chirp when the user presses the button and emit another chirp when the button is released. The users can move freely during the two beeps. He can still correctly select the target device at the direction indicated by the positions of the two beeps. The downside is at the cost of more processing since the joint ETOA detection mechanism cannot be applied since the receivers have no idea about the interval between the two beeps. The selecting device may also need to record and detect the ETOA to set up the ETOA filter to filter out faked ETOA attacks.

ETOA exchange over acoustic channel One prerequisite condition for P&C to work is the ability to exchange ETOA messages before the network is set up. Besides the means presented in Section 5.2, acoustic channel can be used as well. The acoustic channel is inherently a broadcast channel and can be established as long as both party have a microphone and a speaker. There are many possible choices of audio codecs that can fulfill the transmission tasks [15]. However, such codecs are sensitive to environmental noises and is usually effective in a short range (up to a few meters).

To enhance the error resilience, and hence the working range, we can further exploit the high accuracy and robust detection of chirp signals, which is proven in the device selection process and in [20]. We term this method as beep interval coding (BIC). Specifically, we can use the interval between consecutive fixed length chirps to modulate the bits. The ETOA (in hex representation) and the MAC address will translate into a series of beeps separated by intervals corresponding to the hex digits of the digest. BIC is doomed to

have low bit rate since most of the air time is occupied by the transmission of chirps. However, BIC is still a valid design because we only need to transmit the ETOA measurement (20 bits, which suffices in practice) and MAC address (32 bits).

Supporting multiple device pairing We do not believe that the case, where multiple pairs of devices are performing simultaneous selections in close proximity, is very common. Note that this case is different from the scenario where multiple devices are performing device pairing at different locations, but only one pair at each location. Our solution works well in the latter case. To handle simultaneous pairing at a single location, we simply defer its pairing operation when a device detects an ongoing pairing procedure. Anyway, each pairing operation only lasts for a few seconds, and this short delay will not affect most users' tasks. Each device needs to listen to the channel for a period of time, say, a few seconds, to ensure that no chirp signal is detected.

Service trigger and energy efficiency Ideally, P&C should run as a background service and on-demand triggered by a pairing request from other devices. This is viable for fixed computing devices such as a desktop and printer that have an external power supply. However, it is not affordable on battery-operating mobile devices because periodical scanning is needed to discover any pairing request. One possible way to amend this is to add minor additional hardware such as an integral circuit for audio sound detection so that the pairing requester can use a high volume sound to trigger the P&C service. But this will significantly limit the applicability.

We have chosen to ask users to turn on the P&C service explicitly. One may think that the *point* and *select* step can be bypassed by letting the target device directly broadcast its MAC address using BIC after P&C service is triggered. However, this is doable only if the users are sure there are no other concurrent sessions going on. Otherwise, the membership of the pairing cannot be guaranteed, and will degrade to trial and error based on visual feedback. The general P&C work flow will not have such potential risks, due to its selection process. And more importantly, the pointing action provides users with an intuitive way to select the target device and ensures a natural user experience.

10. RELATED WORK

We now compare P&C and related work. Many solutions to device pairing have appeared in literature [11–13, 16, 17, 22]. In general, they can be classified into two categories: those requiring extra hardware or infrastructure support, and those working on the standard mobile phone.

In the first category, extra sensors/hardware (e.g., accelerator, visual markers, Ired or NFC) or infrastructure support will be needed to identify which pair of devices are to be connected. Smart-Its-Friends scheme [12] uses motion sensors to capture the movement patterns when the pairing devices are placed and moved together. The recorded sensing data will be used to identify the target later on. Other solutions such as synchronous gestures [11], Shake Well Before Use [16], and Martini Synch [13] use sensors to record synchronous actions by users (e.g., device shaking [16]) or events on both devices (e.g., one bump into the other [11]), so that such sync information can be used to detect the right pairing device. Toss-It [31] uses accelerometers to detect user motion (toss and swing) for device pairing. Other approaches rely on localization infrastructure to infer spatial reference, in order to differentiate the target from other devices, for example, [17] uses both radio frequency and ultrasonic communications to measure the relative position of the pairing device. Other work also adopts the pointing action to

select pairing devices, including those using a laser pointer or the camera with 3D visual processing [19, 27], or visual marks like QR code [2] and Semacode [3]. [24] compares their performance and usability to identify devices. P&C is different from these solutions in that it does not use any extra hardware or infrastructure, and still preserves excellent usability.

The second category of solution works on standard COTS phones. One prominent example is the “scan-browse-select” scheme used by Bluetooth pairing. Each candidate device is identified with a name or address, and the selecting device scans all the nearby devices and retrieves the name/address list. The user then browses the list and selects the one she believes as the right one. Another proposal SyncTap [22] requires both users to simultaneously press and release the buttons on both devices. It still uses the synchronous action to achieve device pairing but does not require extra sensors. Amigo [30] uses the common radio profile specific to a given location and time to differentiate the pair of devices. It works well when the two pairing devices are in close physical proximity but far away from others. However, the physical proximity identified using radio information is of coarse granularity, and it is hard to identify the target device if it is surrounded by others. P&C is different from these solutions in that the operation is simpler, and does not require explicit user synchronization from both people.

There is also significant research effort focusing on the security aspect of device pairing. They establish a shared secret (usually a temporary key or PIN) to differentiate the device pair from others, through three popular techniques: physical contact, PIN input, and an out-of-band (OOB) channel. In the physical contact based approach, several industry proposals exist, including Wireless USB Association Models [6] using USB cables, WiFi Protected Setup [4] with its practice Windows Connect Now [5] using Ethernet or USB cable and Bluetooth Simple Pairing [1] using Near Field Communication (NFC). They all use at least one type of auxiliary channel to connect both pairing parties. In the PIN based approach, users may need to input keys on the pairing devices. The common practice is a standard Bluetooth pairing scheme. UIA [9] uses a more user-friendly key introduction, consisting of three words randomly chosen from a dictionary. Users need to tell each other their keys and select accordingly from screen prompts. [29] compares usability of simple PIN-based pairing schemes. In the OOB based approach, infrared, visual and audio channels are used to establish authentication and secrecy. These OOB channels, formed through the sensor and actuator pairs available on the devices, are used to verify whether the keys computed at both devices are identical. Talking-to-strangers [8] illustrates a solution based on bidirectional infrared channels or other location-limited channels. Seeing-Is-Believing [18] and ViC [25] utilize visual channels consisting of barcodes or a blinking LED for pairing; Loud-and-Clear [10] and HAPADEP [28] devise schemes using audible channels with visual assistance. [21] and [26] use auxiliary visual and audio channels to compare short and simple synchronized audiovisual patterns (i.e., “beeping” and “blinking”). P&C does not focus on the authentication protocol in device pairing, but addresses a few practical attacks. Therefore, P&C complements these authentication schemes and may use one of them, e.g., the OOB-based approach, to further ensure device-level authentication.

11. CONCLUSION

In this research, we have designed, implemented, and evaluated Point&Connect, an intention-based device pairing solution for mobile phone users. To pair one's phone with another device, a user simply points her phone towards the intended target, in a setting where many other devices may be present in the locality. Each

device captures the pointing gesture by measuring the distance reduction via acoustic sensing techniques. The user then selects the device that observes the largest distance change during the gesture. Our experiments have confirmed the effectiveness of the solution.

In a broader context, P&C explores a new design paradigm for solutions over COTS phones. It seeks to exploit simple user action and perception, combined with standard mobile device capabilities, to achieve features that were previously available only to high-end phones with additional sensors or with infrastructure support. P&C relies on human perception capabilities, which serve as *sensors* and *actuators* for the device, to offset the limitations of COTS phones. Through our experience of P&C, we show that it is indeed a cool and effective way to achieve intuitive device pairing.

12. ACKNOWLEDGEMENT

We would like to thank Mr. Dongkun Kim (Yonsei Univ., Korea) and Bruce Yang (MSR Asia) for their support in the initial implementation of the system. We also thank those research interns at MSR Asia for participating our user study. Our particular thanks go to all the anonymous reviewers for their constructive, technical and editorial comments and advices, and Dr. Roy Want for shepherding the final version of the paper.

13. REFERENCES

- [1] Bluetooth simple pairing.
http://www.bluetooth.com/NR/rdonlyres/0A0B3F36-D15F-4470-85A6-F2CCFA26F70F/0/SimplePairing_WP_V10r00.pdf.
- [2] QR code. <http://www.nttdocomo.co.jp/english/service/imode/make/content/barcode/tool/>.
- [3] Semacode. <http://semacode.com/>.
- [4] WiFi Protected Setup.
<http://www.wi-fi.org/wifi-protected-setup>.
- [5] Windows Connect Now Technology.
<http://www.microsoft.com/windowsxp/using/networking/getstarted/windowsconnectnow.mspx>.
- [6] Wireless USB Association Models.
<http://www.usb.org/developers/wusb/>.
- [7] A. Adya, P. Bahl, R. Chandra, and L. Qiu. Architecture and techniques for diagnosing faults in IEEE 802.11 infrastructure networks. In *MobiCom '04: Proceedings of the 10th annual international conference on Mobile computing and networking*, pages 30–44, New York, NY, USA, 2004. ACM.
- [8] D. Balfanz, D. K. Smetters, P. Stewart, and H. C. Wong. Talking to strangers: Authentication in ad-hoc wireless networks. In *Proceedings of Network and Distributed System Security Symposium 2002 (NDSS'02)*, San Diego, CA, February 2002.
- [9] B. Ford, J. Strauss, C. Lesniewski-Laas, S. Rhea, F. Kaashoek, and R. Morris. Persistent personal names for globally connected mobile devices. In *OSDI '06: Proceedings of the 7th symposium on Operating systems design and implementation*, pages 233–248, Berkeley, CA, USA, 2006. USENIX Association.
- [10] M. T. Goodrich, M. Sirivianos, J. Solis, G. Tsudik, and E. Uzun. Loud and clear: Human-verifiable authentication based on audio. In *ICDCS '06: Proceedings of the 26th IEEE International Conference on Distributed Computing Systems*, page 10, Washington, DC, USA, 2006. IEEE Computer Society.
- [11] K. Hinckley. Synchronous gestures for multiple persons and computers. In *UIST '03: Proceedings of the 16th annual ACM symposium on User interface software and technology*, pages 149–158, New York, NY, USA, 2003. ACM.
- [12] L. E. Holmquist, F. Mattern, B. Schiele, P. Alahuhta, M. Beigl, and H.-W. Gellersen. Smart-its friends: A technique for users to easily establish connections between smart artefacts. In *UbiComp '01: Proceedings of the 3rd international conference on Ubiquitous Computing*, pages 116–122, London, UK, 2001. Springer-Verlag.
- [13] D. Kirovski, M. Sinclair, and D. Wilson. The martini synch: Device pairing via joint quantization. In *ISIT '07: Proceedings of IEEE International Symposium on Information Theory, 2007*, July 2007.
- [14] B. Kusy, A. Ledeczi, and X. Koutsoukos. Tracking mobile nodes using rf doppler shifts. In *SenSys '07: Proc. of the 5th Intl. conference on Embedded networked sensor systems*, pages 29–42, New York, NY, USA, 2007. ACM.
- [15] C. V. Lopes and P. M. Q. Aguiar. Aerial Acoustic Communications. In *Proceedings of the 2001 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, New Paltz, NY, 2001.
- [16] R. Mayrhofer and H. Gellersen. Shake well before use: Authentication based on accelerometer data. In A. LaMarca, M. Langheinrich, and K. N. Truong, editors, *Pervasive*, volume 4480 of *Lecture Notes in Computer Science*, pages 144–161. Springer, 2007.
- [17] R. Mayrhofer, H. Gellersen, and M. Hazas. Security by spatial reference: Using relative positioning to authenticate devices for spontaneous interaction. In *UbiComp*, pages 199–216, 2007.
- [18] J. M. McCune, A. Perrig, and M. K. Reiter. Seeing-is-believing: Using camera phones for human-verifiable authentication. In *SP '05: Proceedings of the 2005 IEEE Symposium on Security and Privacy*, pages 110–124, Washington, DC, USA, 2005. IEEE Computer Society.
- [19] K. Nickel and R. Stiefelhagen. Visual recognition of pointing gestures for human-robot interaction. *Image Vision Comput.*, 25(12):1875–1884, 2007.
- [20] C. Peng, G. Shen, Y. Zhang, Y. Li, and K. Tan. Beepbeep: a high accuracy acoustic ranging system using cots mobile devices. In *SenSys '07: Proceedings of the 5th international conference on Embedded networked sensor systems*, pages 1–14, New York, NY, USA, 2007. ACM.
- [21] R. Prasad and N. Saxena. Efficient device pairing using “human-comparable” synchronized audiovisual patterns. In S. M. Bellovin, R. Gennaro, A. D. Keromytis, and M. Yung, editors, *ACNS*, volume 5037 of *Lecture Notes in Computer Science*, pages 328–345, 2008.
- [22] J. Rekimoto, Y. Ayatsuka, and M. Kohno. Synctap: An interaction technique for mobile networking. In *Mobile HCI*, pages 104–115, 2003.
- [23] E. Rukzio. *Physical Mobile Interactions: Mobile Devices as Pervasive Mediators for Interactions with the Real World*. PhD thesis, University of Munich, 2006.
- [24] E. Rukzio, K. Leichtenstern, V. Callaghan, A. Schmidt, P. Holleis, and J. Chin. An experimental comparison of physical mobile interaction techniques: Touching, pointing and scanning. In *UbiComp'06: Eighth International Conference on Ubiquitous Computing*, September 2006.
- [25] N. Saxena, J.-E. Ekberg, K. Kostiaainen, and N. Asokan. Secure device pairing based on a visual channel. In *SP '06: Proceedings of the 2006 IEEE Symposium on Security and Privacy*, pages 306–313, Washington, DC, USA, 2006. IEEE Computer Society.
- [26] N. Saxena, M. B. Uddin, and J. Voris. Universal device pairing using an auxiliary device. In *SOUPS '08: Proceedings of the 4th symposium on Usable privacy and security*, pages 56–67, New York, NY, USA, 2008. ACM.
- [27] G. Schmidt, Y. Baillot, D. G. Brown, E. B. Tomlin, and J. E. I. Swan. Toward disambiguating multiple selections for frustum-based pointing. In *3DUI '06: Proceedings of the 3D User Interfaces*, pages 87–94, Washington, DC, USA, 2006. IEEE Computer Society.
- [28] C. Soriente, G. Tsudik, and E. Uzun. Hapadep: Human-assisted pure audio device pairing. In *ISC*, pages 385–400, 2008.
<http://eprint.iacr.org/2007/093.pdf>.
- [29] E. Uzun, K. Karvonen, and N. Asokan. Usability analysis of secure pairing methods. <http://research.nokia.com/files/NRC-TR-2007-002.pdf>, 2007.
- [30] A. Varshavsky, A. Scannell, A. LaMarca, and E. de Lara. Amigo: Proximity-based authentication of mobile devices. In *UbiComp*, pages 253–270, 2007.
- [31] K. Yatani, K. Tamura, K. Hiroki, M. Sugimoto, and H. Hashizume. Toss-it: intuitive information transfer techniques for mobile devices. In *CHI Extended Abstracts*, pages 1881–1884, 2005.