

Using IMUs to Identify Supervisors on Touch Devices

Ahmed Kharrufa¹(✉), James Nicholson¹, Paul Dunphy¹,
Steve Hodges², Pam Briggs³, and Patrick Olivier¹

¹ Culture Lab, Newcastle University, Newcastle upon Tyne, UK
{ahmed.kharrufa, james.nicholson, paul.dunphy,
patrick.olivier}@ncl.ac.uk

² Microsoft Research, Cambridge, UK
steve.hodges@microsoft.com

³ PaCT Lab, Northumbria University, Newcastle upon Tyne, UK
p.briggs@northumbria.ac.uk

Abstract. In addition to their popularity as personal devices, tablets, are becoming increasingly prevalent in work and public settings. In many of these application domains a supervisor user – such as the teacher in a classroom – oversees the function of one or more devices. Access to supervisory functions is typically controlled through the use of a passcode, but experience shows that keeping this passcode secret can be problematic. We introduce SwipeID, a method of identifying supervisor users across a set of touch-based devices by correlating data from a wrist-worn inertial measurement unit (IMU) and a corresponding touchscreen interaction. This approach naturally supports access at the time and point of contact and does not require any additional hardware on the client devices. We describe the design of our system and the challenge-response protocols we have considered. We then present an evaluation study to demonstrate feasibility. Finally we highlight the potential for our scheme to extend to different application domains and input devices.

Keywords: IMU · Association · Authentication · Touch interaction · UI design

1 Introduction

Touch-based computing devices, and in particular a variety of tablet form factors, are becoming prevalent. Initially used as personal devices, they are increasingly being used in work and public settings such as schools, shops, museums and exhibition spaces. The focus in this paper will be on school settings, as a principal example, considering that tablet-based classroom applications are growing rapidly [3] and are seen by many as one of the key classroom innovations of the 21st century. However, the same approach applies to the other settings and for other touch devices such as interactive boards, tabletops, laptops and other devices with interactive screens.

In order to leverage the full benefits that tablets can offer in a classroom setting, a teacher or classroom assistant will often need to override the settings on a student's

device and initiate machine-level or even classroom-level effects such as projecting the work of one student onto the class whiteboard or gaining temporary Internet access [2, 8, 13, 14, 19]. Imagine the following scenario: a student has shown an interesting




Fig. 1. SwipeID. Associating touch data with IMU data using a challenge-response protocol.

approach to solving a problem which the teacher wants to share with other students. The teacher initiates supervisor access on that student's tablet, then freezes all the other devices in the classroom and projects to the classroom display for a class discussion. Moving through the class, the teacher and one or more teaching assistants can then authenticate other students' devices to pull content from the board, or to sanction sharing of content between students.

A key requirement to support scenarios like these is the ability to identify and authenticate interactions by supervisory users such as teachers, and to differentiate these from regular users, in this case the students. Here we encounter a known problem – secure authentication is difficult to achieve with any tablet or surface interface, given the ease with which casual observers can engage in “shoulder surfing”. Most commonly, user authentication is based on “something you know” that is meant to be held secret – such as a password or PIN. However, this passcode secret is overly exposed when using touch screen input [15] and can be easily compromised in a public setting such as a classroom [15] where authentication is typically in full view of a number of observers. This inevitably leads to uncontrolled student access to unauthorized applications.

Controlling student access to privileged functions is a very real problem. For example, in 2013 Apple was awarded a \$50 million contract from the Los Angeles School Board of Education, to roll iPads out into public schools across the state – intended to be the first of many such large scale education initiatives. However a year later, the scheme faltered in large part because of identity management and authentication concerns raised when the students found ways to access unauthorized content and applications, leading to significant problems with classroom discipline and ultimately challenging broadband capacity [17, 23, 25].

New authentication and access management solutions in the form of biometrics and near field communication (NFC) may help, but these are only available in some of the latest tablets and there is a pressing need for effective, usable access control for existing

tablets including the estimated 10 million iPads *already* in the classroom in the US alone [17].

In this paper, we offer a novel and elegant authentication solution called SwipeID that simplifies supervisory access to tablets with no NFC or biometric capability. Importantly, it requires no hardware augmentation of the touch device itself. Instead, our solution comprises a wrist mounted inertial measurement unit (IMU) worn only by the supervisor(s), and a challenge-response style interaction protocol involving a very simple set of movements. This setup is depicted in Fig. 1. The aim is to provide a practical alternative when other solutions such as NFC and fingerprint readers are not possible (e.g. for iPads and interactive boards in schools, and for large/fixed touch displays in public settings that do not normally have NFC, fingerprint readers, or cameras). Unlike other options, ours allows for identification at the point of touch – especially useful for large displays like interactive boards. We precede a more detailed description and evaluation of our solution with an overview of related work in the areas of touch and sensor interactions and their deployment in a classroom setting.

Our contribution is as follows: We present a novel system that allows simple, point of contact authentication for any touch screen device. We show how such a system can be used to solve known access management problems when using tablets and other touch screens in the classroom as an example. We demonstrate the efficacy of our system in user studies with four exemplar user interface controls.

2 Related Work

In this section we cover previous work which tackles the challenge of identifying users of touch devices. We also briefly review literature relating to the use of movement correlations for building trust, a very relevant topic. Finally we consider previous approaches to the specific challenge of teacher orchestration in the classroom, a key application area which inspired this work.

2.1 Identifying Touches

Much work that seeks to identify users on touch devices has focused on the tabletop context. In DiamondTouch [7] sensors are embedded in the chairs of users. In conjunction with a transmitter built into the display it is possible to identify the owner of each touch event. Roth et al. [28] proposed the *IR Ring*; an infra-red (IR) augmented finger-worn ring that transmits IR pulses detectable by cameras embedded into a tabletop. Particular patterns of pulses are unique to particular users, making them suitable for identifying touches across sessions and across different tabletops. A similar approach using a wristband was proposed by Meyer and Schmidt [21]. However, both technologies are only suitable for IR-sensitive optical touch detection systems.

Holz and Baudisch [12] designed a special touch screen to support biometric user authentication based on fingerprint recognition carried out dynamically during each touch interaction. This provides natural per-touch user identification but relies on high resolution, high speed scanning and processing hardware. Harrison et al. [10]

introduced an approach for user identification that relies on sensing electrical properties of the human body when using capacitive touch devices. However, while promising, the experimental results showed that the variability of these electrical properties due to biological and environmental factors can be larger than the variability of such properties between users. Mock et al. [22], on the other hand, explored using raw sensor data from typing on an interactive display for user identification. The system was based on optical touch sensing rather than the more common capacitive touch displays.

HandsDown [30] and MTi [4] are two approaches that rely on handprints for identification against a database of users' hands characteristics rather than fingerprints. Unlike HandsDown, MTi is not limited to camera-based touch sensing. However, for both the touch surface needs to be big enough to accommodate an outstretched hand. Other approaches (e.g. [18, 26]) use an overhead camera for tracking and identifying users. These are most suitable for large fixed displays, in particular tabletops.

2.2 Inferring Relationships by Sensor Correlation

Researchers have developed a number of techniques to infer the physical relationship between devices being used in conjunction with each other. The use of accelerometers for making associations between multiple devices has been a subject of many investigations [9, 11, 24]. Fujinami and Pirttikangas [9] used the correlation of accelerometer signals in wrist-worn devices and devices embedded in objects to reason about the identity of an object's user. In Smart-Its [11] two objects held together and shaken were associated based on a correlation threshold. Similarly, Patel et al. [24] proposed the use of a shake/pause gesture sequence to pair a mobile device with a public terminal. If the mobile device produces the same shake/pause sequence as that displayed on the terminal, then it is assumed that this is the correct device with which to establish an association. For devices that have already been paired, Chen et al. [6] explored the design space of joint interactions between a smart watch and a smart phone. One of the proposed interactions was to use the accelerometer data from the smart watch to augment the interactions with the phone. Shrirang et al. [31] relied on the correlation between input from a wrist-worn accelerometer and from the keyboard/mouse of a computer terminal to confirm the continuous presence of a user. The user is logged-out in the absence of such correlation.

PhoneTouch [29] used server-side correlation-in-time to allow the use of phones to select targets on an interactive surface using direct touch. One application of associating a detected touch with the phone that caused it, is to use the phones for user identification. However, the need to rely on computer vision to detect phone touches and to distinguish them from finger touches limits its use to vision based touch screens. Moreover, since the system is not specifically designed for user identification, association fails when the system detects more than one touch within the same recognition time frame making it very susceptible to attacks.

2.3 Classroom Orchestration

With the emergence of affordable tablet and tabletop computers, there has been an increased interest in deploying large numbers of single- and multi-user devices in classroom environments [2, 8, 13, 14, 16, 19, 20]. The support of teacher orchestration of the classroom has repeatedly been identified as a key challenge [2, 13, 19]. Teachers are often provided with remote monitoring and control tools and the ability to project the content of one of the devices to a large classroom display. This approach may be facilitated by providing the teacher with a dedicated device [2, 16, 19, 20].

However, confining acts of orchestration to a single, static device [2] fails to recognize the real nature of teaching a class, which involves dynamic engagements with the whole group, sub-groups and individuals [27, 32]. Alternative proposals include the provision of orchestration functionalities on a teacher's hand-held device such as a tablet [19, 20]. While this improvement was reported to be useful by one teacher [20], that same teacher described how having to hold and interact with such a device limited their ability to work with the students directly on their tables. The realities of classroom environments result in a wide range of situations where holding a tablet will restrict the quality of teachers' interaction with students. In recognition of such restrictions and the need for teachers to interact directly with the students' devices, the TinkerLamp project [8] used a 'TinkerKey' tangible object with specific visual markers on it. This was automatically identified by the students' tabletops, to allow teachers to issue special commands. However, TinkerKey relies on optical multi-touch sensing and is therefore incompatible with capacitive tablets.

3 SwipeID

SwipeID is a system we have developed to support supervisors, such as teachers, as they interact with one or more touch devices. We leverage a wrist-worn IMU connected through Bluetooth to a nearby SwipeID server to identify the teacher's interactions (The server can be any available machine that can connect to both the IMU and the other devices and that can perform simple correlation calculations). A key assumption is that only the supervisors in a particular context are wearing IMUs on their wrist. Each of the touch-based client devices runs the SwipeID client service. The client software allows access to commands and tools that can only be successfully activated by a user wearing a pre-configured IMU.

When a touchscreen user command that requires privileged access is executed, a challenge-response protocol is initiated on-screen, requiring the user to perform a particular sequence of gestures. While interacting with this control, all associated touch data is sent over the network to the server. The server, which is continuously reading IMU sensor data from connected wrist devices, compares this sensor data with the touch data transmitted by the client device and then calculates the correlation between them. If the correlation value is above a predefined threshold, a go-ahead is sent back to the client device. If the correlation is low, a 'reject' is sent back to the client. The client device only communicates with the server when privileged access is requested ensuring that the network is not overloaded during normal usage.

SwipeID has the following properties:

- It allows an arbitrary number of people privileged control of any number of touch-based devices.
- It requires no special hardware on the client devices.
- It removes the need to try and keep a password secret and the need to remember passwords.
- It requires a small and relatively low-cost wrist-worn IMU such as a smart watch for each authorized person, along with a networked machine acting as a server.

3.1 Correlating Touch and Sensor Data

During the challenge-response phase, the magnitude of the acceleration of each touch stroke is derived from both IMU and touch data. In the case of the IMU, the data from the on-board accelerometer, gyroscope and magnetometer must be integrated to derive the linear acceleration of the device independent of any confounding movement such as rotation. By using the magnitude of the linear acceleration there is no need to consider the relative orientation of the IMU and touch sensor which could vary with time and depending on how the device is worn. The two data streams are then matched using a correlation-in-time function [9] (Eq. 1) and the resulting probability is thresholded. The touch and IMU data are collected at the same rate, 33 Hz. The data is resampled to account for the intermittent sampling latencies inherent in the devices used. The X and Y touch data is differentiated twice to calculate acceleration. The data is filtered using a 5-sample moving average before and after each derivative calculation, and the magnitude of the acceleration is calculated for both the touch (d1) and linear acceleration (d2) data. The data correlation was calculated as follows:


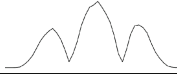



$$r_{12} = \frac{\sum (d_1 - \bar{d}_1) \cdot (d_2 - \bar{d}_2)}{\sqrt{\sum (d_1 - \bar{d}_1)^2 \cdot \sum (d_2 - \bar{d}_2)^2}} \quad (1)$$

3.2 A Network Based Challenge-Response Approach

SwipeID relies on the correlation between the acceleration calculated from the touch data and that measured by the IMU. Accordingly, in theory any user could attempt to trigger a privileged command as authentication can only be carried through performing some gestures aiming to achieve the correlation threshold. This means that without imposing constraints on the type of gestures required for authentication, a user not wearing an IMU could attempt to trigger a privileged command concurrently with a supervisor (who is working on another device), and mirror the movement pattern of the supervisor. If the unauthorized user were to do this well enough, there is a possibility that the correlation calculated between their touch data and that of the IMU on the supervisor's wrist exceeds the threshold and might even be higher than that of the

supervisor - due to measuring acceleration at the wrist which does not map perfectly to the touch data at the fingertip.

Table 1. The magnitude of acceleration from touch data for basic gestures.

Stroke type	Magnitude of touch data acceleration vs time
Single, straight stroke	
A 2 stroke gesture with a 180° change in direction	
A 2 stroke gesture with 90° change in direction	
A 2 stroke gesture with no change in direction, but a small pause between the strokes	
A 2 stroke gesture with 180° change in direction and small pause between the strokes	

To prevent such attacks, a challenge-response approach is used. The server has a pool of distinct challenges equal to or larger than the number of devices in use at any one time. If two users try to gain privileged access at the same time, they will be assigned two different challenges that require different movement patterns. If unauthorized users try to copy the movement pattern of the person wearing the IMU, they will have to deviate from the movement pattern required by their own challenge and thus fail their challenge locally regardless of how well their movement correlates with that of the IMU. Accordingly, the only option users have to gain access to privileged commands is to both accurately follow their own challenges by wearing the IMU.

To best design a set of challenge patterns that ensure distinct movement patterns which will not inadvertently cause high correlation between different challenges, we looked at the basic acceleration signals generated from touch data for the most primitive strokes (Table 1). From these graphs we can see that by interleaving short strokes and periods of no movement, it is possible to generate a number of distinct patterns. If we represent a movement by M and a pause by P, we can design as many different challenges as desired. We want to keep the sequence short to keep it quick to enter, but to decrease the correlation between the different patterns we decided to choose challenges that differ in at least two segments. Table 2 lists 12 different challenges which meet this criteria and also have a minimum of three movements and one pause.

Table 2. Twelve different challenges that combine movements (M) and pauses (P). (See Fig. 2 for example shapes of signals associated with these challenges)

M	P	M	P	M	P
M	P	M	P	P	M
M	P	P	M	M	P
M	P	P	M	P	M
P	M	M	P	M	P
P	M	M	P	P	M
P	M	P	M	M	P
P	M	P	M	P	M
M	M	P	P	M	M
P	M	M	M	M	M
M	M	M	M	M	P
M	M	M	M	P	M

Designing Sensor-Coupled User Controls. To require a user to follow a specific challenge, a custom user interface control must display it and verify the response. We would like these controls to:

1. be user friendly, i.e. feel simple, be quick to use and not too demanding;
2. support the proposed challenge-response protocol; and
3. result in sufficient movement to generate high correlation.


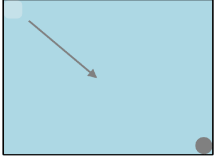


To aim for user friendliness (goal 1), we avoided controls that increase cognitive load, such as requiring users to trace a dynamic path that appears incrementally or changes in real time. We limited the design space to only include controls that simply require users to move along a path or to move to clearly marked targets. With such controls, it is possible to enforce a pattern that results in different levels of movement and to verify that the user is following the displayed challenge (goal 2).

To meet goal 3 we considered two options for imposing distinct changes in the speed of movement:

Present a sequence of targets one at a time. The user is required to move to the next target within a certain time otherwise the challenge fails. Once a target is reached, the subsequent target appears, either immediately or after a certain period, depending on the challenge. This enforces a distinctive movement/no movement pattern. We call this ‘discrete’ interaction as there are explicit wait periods. Note that the user is not aware of the full challenge pattern beforehand.

Show the full challenge, represented by a certain path to navigate through, from the outset. The required pattern of movements is achieved by switching between a wide straight paths which can be navigated through quickly, and shorter narrow and/or curvy paths that require more careful (and slow) navigation. This idea is derived from Accot and Zhai’s work [1] that looked at the relationship between movement time and width, and to a lesser extent curvature, of a path to steer through. We refer to this as ‘continuous’ interaction because no pauses are expected. In this case the users can see the full challenge from the outset.

Table 3. Design space for four different IMU-coupled user interface control types. The shapes in the second row are one of 12 different challenge shapes.

	Small footprint	Large footprint
Discrete Challenge		
Continuous Challenge		




Fig. 2. The magnitude of acceleration signal from touch (black) and IMU (red) for two different challenges using the discrete, small footprint control (Color figure online).

It is hypothesized that users associate different levels of ease of use to options (a) and (b) because of the different style of interaction used. In particular, the use of continuous movement and exposure to the full challenge beforehand is assumed to help in perceiving the interaction as one action (or *chunk* [5]) rather than a number of discrete actions. Pauses or very slow movement in a straight wide path area results in failing the challenge. Moreover, as long as the user does not go outside the path they can speed up the response for a continuous challenge by navigating more quickly. This is not the case with the discrete challenge as the user must wait for each target to appear.

We need to ensure sufficient wrist movement to generate meaningful IMU data. We hypothesized that asking users to perform the task as quickly and as accurately as possible would help with this. We also need to avoid finger-only movements. For this reason, we wanted to compare gestures which use predominantly one axis of movement with those which require movement in two dimensions – with both cases requiring movements above a certain minimum physical distance. The full design space is summarized in Table 3.

For discrete, small footprint interactions, a horizontal path is presented with the next target to move to indicated as a circular region. The following target either appears immediately after reaching this target or after a certain pause, depending on the challenge. Similarly, for the large footprint discrete option, a rectangular shape is displayed with the target appearing at one of the four corners.

For continuous interaction, the control is a path that changes between straight, wide segments and narrow curvy segments that the user needs to navigate (steer) through.




Fig. 3. User study. Discrete, small footprint control (left). Continuous, large footprint control (right).

The path, which ends with a clearly-marked target, is designed to either stretch mostly along one axis (small footprint) or along both (large footprint).

We envisage privileged commands to be presented as normal buttons that expand once touched. When a small discrete challenge button is touched for example, it expands horizontally to the required width showing the next target. The user must then slide their finger without removing it to the next target and so on. Upon sliding to the last point in the challenge, which is the same as the starting point, the button collapses. The button also collapses immediately if the user fails to follow the challenge.

Figure 2 shows the magnitude of acceleration signal at the server side calculated from touch data and that from the IMU for two different challenges. These signals, which are used for correlation calculations, clearly show the similarity between touch and IMU data and also how different challenges result in different signals.

4 User Study

We conducted a study as a proof of concept for SwipeID and to compare the performance and user perception of the four different control types: discrete small footprint, discrete large footprint, continuous small footprint, and continuous large footprint. A repeated measures design was chosen to directly compare participants' performance with each control. We recruited 19 participants for the study (mean age = 26 years; all right handed). Inclusion criteria included good (correctable) vision and some experience with touch-based devices such as smartphones or tablets.

The touch device used was a 10.1" Windows 8 tablet. A commercial IMU, LPMS-B from Life Performance Research was used in the study and was fixed on the

participants' wrists using a sweatband (Fig. 3). This device supports on-board integration of accelerometer, gyro and magnetometer data to calculate linear acceleration. The tablet was placed horizontally on a table and the participants interacted with it from a seated position. All the controls had a physical length of 15.9 cm and the large footprint controls had a height of 12.0 cm.

Initially participants were briefed regarding the purpose of the study and were shown the four control types. They were then given two practice challenges for each control before commencing the study proper. The presentation order for the four controls was counterbalanced and the 12 challenges were randomized for each control (see Table 2 for challenges). Participants were asked to complete each challenge as quickly and accurately as possible. If the participant completed the challenge correctly they were presented with a green feedback screen, whereas a red screen was presented if the challenge was not completed correctly and they were required to retry the challenge until successful. All touch and IMU data from each trial was automatically logged for subsequent analysis.

Upon completing the challenges for each of the four controls, participants were presented with a six-item questionnaire exploring their preferences and perceptions regarding the proposed controls. Four of the questions directly asked the participants to rate their experience using each of the four configurations using a scale ranging from Very Good to Bad. The other questions asked participants to write down their preferred configuration and the configuration they would like to avoid, along with explanations for each selection. Participants were fully debriefed upon completing the questionnaire.

4.1 Results

We performed data analysis to determine the efficacy of our challenge-response protocol. To do this we calculated (i) the correlation between corresponding touch and IMU data and (ii) the similarity between touch data for a challenge performed by the person wearing the IMU and touch data from other challenges.

Table 4 shows the average correlation between touch data and IMU data for the different control types as well as the average time to respond to a challenge and the average number of failed attempts to respond to the challenge out of the total of 12 trials. Failure occurs when the user does not follow the challenge accurately, leading to rejecting the response locally without the need to send data to the server. Upon failure, the user had to repeat the challenge. The table shows that the discrete controls resulted in higher average correlation, shorter average time to completion and a lower number of failed attempts on average than the continuous controls. The receiver operating characteristic (ROC) curves for the four controls (Figs. 4 and 5) show that the discrete controls have better performance than the continuous controls. The ROC curves show that for the discrete controls a threshold of 0.6 allows for 86 % (large footprint) to 88 % (small footprint) true positives and 0.2 % (small footprint) to 0.4 % (large footprint) false positives. As for the continuous controls, while the curve still shows that the controls perform well in separating true positives from false positives, a threshold of 0.5 is probably the best choice, which results in 59 % (small footprint) to 64 % (large footprint) true positives but also allows for 5 % false positives.

Table 4. Mean (and SD) for correlation, execution time and number of failed challenges across the four different control types proposed.

Control type	Mean correlation (SD)	Mean time in msec (SD)	Mean no. of failed challenges (SD)
Discrete, small footprint	0.77 (0.08)	4800 (435)	1.58 (1.77)
Discrete, large footprint	0.71 (0.07)	5280 (362)	1.79 (2.37)
Continuous, small footprint	0.49 (0.11)	7030 (1800)	6.47 (5.47)
Continuous, large footprint	0.53 (0.07)	7130 (1980)	5.53 (3.01)




Fig. 4. Discrete small footprint and discrete large footprint receiver operating characteristic curves.




Fig. 5. Continuous small footprint and continuous large footprint receiver operating characteristic curves.

The results show that participants were faster and failed less challenges when using the discrete small footprint controls, while the mean correlation was also the highest. The discrete large footprint control ranked second overall in the three categories, with continuous small footprint third. The continuous large footprint control performed the worst of all the conditions.

The results also show that the large footprint with horizontal, vertical and diagonal movements did not clearly result in better performance than the small footprint with mostly horizontal movements. In the discrete controls case the small footprint performed marginally better than the large footprint control in terms of correlation, time, and failed challenges. In the continuous controls case the large footprint control performed slightly better in terms of correlation and errors, but worse with regards to time. For the discrete control case, one reason why the large footprint control performed slightly worse than its small counterpart may have been the result of participants having to sometimes lift their wrist to see if the next target had been obscured by their hands. This could have led to unwarranted wrist movements and thus IMU acceleration that did not correspond to any touch movement. Such a scenario would result in a reduced correlation between the touch and the IMU data. For the continuous interaction-large footprint control, participants were tracing a path gradually thus no such wrist movements were observed. The higher average correlation in the continuous interaction-large footprint control compared to its small counterpart may be due to the fact that the horizontal only movement could be performed with more finger movement and less wrist movement as compared to a gesture performed in both horizontal and vertical dimensions.

4.2 Participant Feedback

All the participants were asked to answer a simple questionnaire to provide subjective ‘experience ratings’ regarding the four control types used in the study. Participants were asked to rate their experience of using each of the four controls on a scale from 1 (bad) to 4 (very good). Additionally, participants were asked to select the control they would prefer to use on a daily basis and also indicate the control they most wanted to avoid. Below we present the findings from the questionnaire.

After aggregating the participants’ experience rating scores for each of the four techniques, we found that the ratings for three of the controls were very close. The discrete small footprint came top with a rating of 3.3, while the discrete large footprint and continuous small footprint controls shared a score of 3.2. The continuous large footprint technique received the lowest rating: 2.8.

Despite the discrete small footprint control receiving the highest experience rating, the most popular control amongst participants (the one most participants selected as their favorite for daily use) was the discrete large footprint technique (see Table 5). This control was perceived to be fast whilst yielding a low number of errors, but participants also noted, as we have observed as well, that some angles could be obscured at times by the placement of the hand (e.g. when target is located in lower-right hand corner) and waiting for the dot to move could become tedious. This may help explain why, despite being the preferred technique by some, it also had a

Table 5. Summary of participants' control preference ratings

Control	Preferred choice (%)	Avoid (%)
Discrete small footprint	31.6	10.5
Discrete large footprint	36.8	21.1
Continuous small footprint	31.6	10.5
Continuous large footprint	0	57.9

higher avoidance percentage than both the discrete and continuous small footprint controls by others. The general consensus regarding the discrete small footprint control, the tied-second most popular technique, was that this technique was fast, easy and intuitive.

With regards to the continuous controls, participants agreed that the small footprint technique was fast, but some argued that it was potentially too much work for identification. However, while the small footprint control shared almost the same preference rating as both types of discrete challenge controls, the large footprint version did not obtain any preference votes, and received the majority of avoidance votes. Users mentioned that the technique required too much concentration (i.e., was too much work) and caused too many errors resulting in a 'frustrating' experience. This was a somewhat surprising considering that it had slightly lower failure rate and slightly higher average correlation than the small footprint version, although the average completion time was slightly longer.

The results of our questionnaire show that despite quantitatively being the fastest, leading to fewer errors, and being the highest rated control in terms of experience, the discrete small footprint control was only the second most popular choice amongst participants. However, it is unclear why this was the case as there was no reported negative feedback on this technique. It may be a case of participants' perception of the large footprint being more suitable for the identification task. Clearly, however, participants did not favor the continuous large footprint technique.

5 Discussion and Future Work

SwipeID can identify any user wearing the IMU on any touch-based device without the need for special hardware on the touch device, and at the point of contact on the screen. This approach can be a practical solution for devices that do not have a fingerprint or NFC reader and, most immediately, we envision it as a practical solution for supervisor authentication to iPads already in classroom use. Accordingly, we view the decision of using SwipeID as one of practicality rather than based on performance measures alone when compared to NFC, biometric, or password authentication. For example, with devices that do not have NFC or biometric, password is the most likely alternative but, as we have discussed earlier, passwords are not as practical with touch devices in public settings as they are more prone to shoulder surfing especially when it is not possible to use the device in a private environment. Moreover, even with devices that do have NFC or biometric identification, SwipeID could be useful in scenarios where identification is required at a specific point on a large screen where it may be

impractical to move to a specific location for biometric or NFC identification. This includes collaborative work on a large classroom whiteboard where passcode-based authentication procedures are even more visible and could be easily compromised, or when collaborating around digital tabletops for example where issues of reach caused by the large surface render other techniques impractical. With SwipeID authentication is done at the point of contact using simple gestures. The combined use of a challenge-response protocol and our special user control ensure that only the user wearing the IMU is able to successfully respond to the challenge. While this approach does have a certain level of technical complexity, from users' perspective it is a simple technique that is unobtrusive and makes only few assumptions on the user's part.

One of the motivations for this work was to improve the ways in which teachers can interact with students and their touch-based devices in a classroom setting. We note that SwipeID offers a number of opportunities in this space, given that it allows for freedom of movement around the class and allows for device or screen-specific authentication at the point of contact. It also removes the need for special hardware on the students' devices, the use of a secret passcode, or having to issue commands through a separate fixed or handheld computer. As we noted in the introduction, tablets are increasingly used in the classroom, but their uptake is currently being limited by known access management concerns. Better tools for user identification are therefore an important educational issue. However, considering that even with discrete controls, authentication gestures took 4–5 s, this may be perceived to be too long in a classroom context. Exploring the use of SwipeID in a real classroom setting and exploring options for reducing the required gesture time for authentication are two important areas for future investigation.

SwipeID can also be used in other scenarios where supervisor access of touch-based interactions needs to be identified. For example, in a retail context a sales assistant may need to configure the information displayed in-store, while customers are only allowed to browse the information. Alternatively, in a museum a staff member might need to dynamically update information displayed on interactive terminals. Moreover, SwipeID lends itself well to devices with large or fixed interactive displays where it is not possible to interact with the display privately without being prone to shoulder surfing.

In some applications, multiple users share touch-based devices or need to interact collaboratively on a large interactive surface. In these cases it can be beneficial to identify each specific individual in a way that is resistant to shoulder-surfing – again, a known problem with traditional authentication on touch screen devices. If each user is able to use a wrist-worn motion sensor then SwipeID can be used to identify them.

Another possibility for future work is to explore the use of SwipeID with mouse interactions for user authentication (rather than just verifying user presence as in Shrirang et al. [31]). This could allow the use of a single system as a universal authentication solution across a full range of devices within a particular context – educational or otherwise. However we should note that a key challenge with using a mouse is that the user can perform relatively large gestures on the screen without a significant movement of the wrist. The level of correlation between the physical movement of the mouse and the movement of the mouse pointer is also highly dependent on the gain setting of the mouse so this would also need to be factored in.

Our user study aimed to demonstrate the validity of our approach and to evaluate the performance and perception of four proposed controls. It showed that discrete controls performed better than continuous controls in all regards. However, we observed that when users failed a challenge and had to repeat it for the continuous controls case, they performed the gesture faster which resulted in better correlation for that challenge. This appears to support the theory that it is possible for users to improve their performance over time with continuous controls, giving them a longer term advantage over the discrete controls. In future work it may be possible to further optimize the continuous controls to improve performance.

A limitation of our study was that it was conducted in a lab environment rather than in one of the contexts within which we claim its utility. Future work needs to investigate its longitudinal use in an ‘in-the-wild’ environment where supervisors will be interacting with devices of different form factors and from different seating and standing positions. A longitudinal, in the wild evaluation will also help in gaining better understanding of the user’s preferences and may explain some of the discrepancies between preferences and performances recorded in Tables 4 and 5. It can also show whether the measured failure rate can improve with repeated use and whether the level it settles at can cause annoyance to users or not.

SwipeID user identification is reliant on a wrist worn IMU. In other words, the framework identifies the IMU and not the actual user. This means that whoever has the IMU will have privileged access regardless of the actual identity of the user (relying on what the user has rather than what the user knows). This means that the main threat for the system is getting access to the IMU by unauthorized users. This threat is common with other techniques that rely on a hardware key for authentication such as NFCs. Designing the system to overcome this limitation was outside the scope of the work reported here, but one possible solution would be to require a secret passcode to be entered into the wrist-worn IMU prior to use. The user would only be required to enter the passcode when the wrist-band with the IMU is first put on, which will normally just be once and could be done discreetly (thus adding the requirement of something the user knows as well). An alternative to the passcode is to take advantage of the motion sensors and use behavioral biometric to continuously authenticate the person wearing the sensor.

With the increasing popularity of wrist-worn devices incorporating motion sensors, such as smart watches, the need for a dedicated wrist-worn sensor could ultimately be eliminated. While our current work uses linear acceleration data from the IMU which is derived from data collected from an accelerometer, a gyroscope and a magnetometer, future work will investigate the use of accelerometer data only, thus eliminating the need for the gyroscope and the magnetometer, reducing cost and increasing compatibility with wrist-worn consumer devices.

6 Conclusion

In this paper we considered the problem of identifying supervisors across multi-touch devices and gave an illustrative classroom scenario. We proposed SwipeID, a system that works across any touch device and which provides an IMU to supervisors and

searches for correlations between the IMU and touch events on known devices to identify the display being used by a supervisor. In our user study we explored IMU-coupled user interface controls to accompany the proposed system and found that requiring the user to perform discrete gestures, on small interface controls enabled us to identify the best correlation between touch and IMU data. However, the time required to complete the authentication process may be perceived to be too long in certain contexts, thus reducing authentication time is identified as an area for further investigation. We propose that SwipeID can also be used across multiple devices, that it has significant classroom potential and that future work can explore its applicability with other input techniques and its relevance to other contexts.

Acknowledgments. This work was supported by the RCUK Digital Economy Programme-SIDE: Social Inclusion through the Digital Economy EP/G066019/1

References

1. Accot, J., Zhai, S.: Beyond Fitts' law: models for trajectory-based HCI tasks. In: Proceedings CHI 1997, pp. 295–302. ACM Press, New York (1997)
2. AlAgha, I., Hatch, A., Ma, L., Burd, L.: Towards a teacher-centric approach for multi-touch surfaces in classrooms. In: Proceedings ITS 2010, pp. 187–196. ACM, New York (2010)
3. British Educational Suppliers Association (BESA): Tablets and apps in schools 2013. ICT Series, May 2013
4. Blažica, B., Vladušič, D., Mladenčić, D.: MTi: a method for user identification for multitouch displays. *Int. J. Hum. Comput. Stud.* **71**(6), 691–702 (2013)
5. Buxton, W.A.: Chunking and phrasing and the design of human-computer dialogues. In: Baecker, R.M., Grudin, J., Buxton, W.A., Greenberg, S. (eds.) *Human-Computer Interaction: Toward the Year 2000*, pp. 494–499. Morgan Kaufmann Publishers, San Francisco (1995)
6. Chen, X.A., Grossman, T., Wigdor, D.J., Fitzmaurice, G.: Duet: exploring joint interactions on a smart phone and a smart watch. In: Proceedings CHI 2014, pp. 159–168. ACM Press, New York (2014)
7. Dietz, P., Leigh, D.: DiamondTouch: a multi-user touch technology. Mitsubishi technical report (2003)
8. Do-Lenh, S.: Supporting reflection and classroom orchestration with tangible tabletops. Ph.D thesis, École Polytechnique Fédérale de Lausanne (2012)
9. Fujinami, K., Pirttikangas, S.: A study on a correlation coefficient to associate an object with its user. In: Proceedings 3rd IET International Conference on Intelligent Environments (IE 2007), pp. 288–295 (2007)
10. Harrison, C., Sato, M., Poupyrev, I.: Capacitive fingerprinting: exploring user differentiation by sensing electrical properties of the human body. In: Proceedings UIST 2012, pp. 537–544. ACM Press, New York (2012)
11. Holmquist, L.E., Mattern, F., Schiele, B., Alahuhta, P., Beigl, M., Gellersen, H.W.: Smart-its friends: a technique for users to easily establish connections between smart artefacts. In: Proceedings Ubicomp 2001, pp. 116–122. ACM Press, New York (2001)
12. Holz, C., Baudisch, P.: Fiberio: a touchscreen that senses fingerprints. In: Proceedings UIST 2013, pp. 41–50. ACM Press, New York (2013)

13. Kharrufa, A., Balaam, M., Heslop, P., Leat, D., Dolan, P., Olivier, P.: Tables in the wild: lessons learned from a large-scale multi-tabletop deployment. In: Proceedings CHI 2013, pp. 1021–1030, ACM Press (2013)
14. Kharrufa, A., Martinez-Maldonado, A., Kay, J., Olivier, P.: Extending tabletop application design to the classroom. In: Proceedings ITS 2013, pp. 115–124. ACM, New York (2013)
15. Kim, D., Dunphy, P., Briggs, P., Hook, J., Nicholson, J.W., Nicholson, J., Olivier, P.: Multi-touch authentication on tabletops. In: Proceedings CHI 2010, pp. 1093–1102. ACM Press, New York (2010)
16. Kreitmayer, S., Rogers, Y., Laney, R., Peake, S.: UniPad: orchestrating collaborative activities through shared tablets and an integrated wall display. In: Proceedings UbiComp 2013, pp. 801–810. ACM Press, New York (2013)
17. Leonard, D.: The iPad goes to school (2013). <http://www.businessweek.com/articles/2013-10-24/the-ipad-goes-to-school-the-rise-of-educational-tablets>
18. Martinez, R., Collins, A., Kay, J., Yacef, K.: Who did what? Who said that? Collaid: an environment for capturing traces of collaborative learning at the tabletop. In: Proceedings ITS 2011, pp. 172–181. ACM Press, New York (2011)
19. Martinez-Maldonado, R., Kay, J., Yacef, K., Edbauer, M., Dimitriadis, Y.: MTClassroom and MTDashboard: supporting analysis of teacher attention in an orchestrated multi-tabletop classroom. In: Proceedings CSCL 2013, pp. 320–327 (2013)
20. Mercier, E., McNaughton, J., Higgins, S., Burd, E., Maldonado, R.M., Clayphan, A.: Interactive surfaces and spaces: a learning sciences agenda. In: ICLS 2012 (2012)
21. Meyer, T., Schmidt, D.: IdWristbands: IR-based user identification on multi-touch surfaces. In: Proceedings ITS 2010, pp. 277–278. ACM Press, New York (2010)
22. Mock, P., Edelmann, J., Schilling, A., Rosenstiel, W.: User identification using raw sensor data from typing on interactive displays. In: Proceedings Intelligent User Interfaces (IUI 2014), pp. 67–72. ACM Press, New York (2014)
23. Needle, D.: L.A. school district puts the brakes on massive iPad deployment (2014). <http://tabtimes.com/feature/ittech-os-ipad-ios/2014/08/26/la-school-district-puts-brakes-massive-ipad-deployment>
24. Patel, S.N., Pierce, J.S., Abowd, G.D.: A gesture-based authentication scheme for untrusted public terminals. In: Proceedings UIST 2004, pp. 157–160. ACM Press, New York (2004)
25. Pierra, P.: Tablets-in-school galore causes a bandwidth war in Southern California (2014). <http://tabtimes.com/case-study/education/2014/05/15/tablets-school-galore-causes-bandwidth-war-southern-california>
26. Ramakers, R., Vanacken, D., Luyten, K., Coninx, K., Schöning, J.: Carpus: a non-intrusive user identification technique for interactive surfaces. In: Proceedings UIST 2012, pp. 35–44. ACM Press, New York (2012)
27. Reid, D.J.: Spatial involvement and teacher-pupil interaction patterns in school biology laboratories. *Educ. Stud.* **6**(1), 31–41 (1980)
28. Roth, V., Schmidt, P., Güldenring, B.: The IR ring: authenticating users' touches on a multi-touch display. In: Proceedings UIST 2010, pp. 259–262. ACM Press, New York (2010)
29. Schmidt, D., Chehimi, F., Rukzio, E., Gellersen, H.: PhoneTouch: a technique for direct phone interaction on surfaces. In: Proceedings UIST 2010. pp. 13–16. ACM Press, New York (2010)
30. Schmidt, D., Chong, M.K., Gellersen, H.: HandsDown: hand-contour-based user identification for interactive surfaces. In: Proceedings of the 6th Nordic Conference on Human-Computer Interaction: Extending Boundaries, pp. 432–441 (2010)

31. Mare, S., Markham, A.M., Cornelius, C., Peterson, R., Kotz, D.: ZEBRA: zero-effort bilateral recurring authentication. In: Proceedings of the 2014 IEEE Symposium on Security and Privacy (SP 2014), pp. 705–720. IEEE Computer Society, Washington, DC (2014)
32. Smith, H.A.: Nonverbal communication in teaching. *Rev. Educ. Res.* **49**(4), 631–672 (1979)