A Pilot Study on the Security of Pattern Screen-Lock Methods and Soft Side Channel Attacks

Panagiotis Andriotis*
University of Bristol, MVB
Bristol, BS8 1UB, UK
p.andriotis@bristol.ac.uk

Theo Tryfonas
University of Bristol, QB
Bristol, BS8 1TR, UK
theo.tryfonas@bristol.ac.uk

George Oikonomou
University of Bristol, MVB
Bristol, BS8 1UB, UK
g.oikonomou@bristol.ac.uk

Can Yildiz
University of Bristol, UK
canyildiz.2011@my.bristol.ac.uk

ABSTRACT

Graphical passwords that allow a user to unlock a smartphone's screen are one of the Android operating system's features and many users prefer them instead of traditional text-based codes. A variety of attacks has been proposed against this mechanism, of which notable are methods that recover the lock patterns using the oily residues left on screens when people move their fingers to reproduce the unlock code. In this paper we present a pilot study on user habits when setting a pattern lock and on their perceptions regarding what constitutes a secure pattern. We use our survey's results to establish a scheme, which combines a behaviour-based attack and a physical attack on graphical lock screen methods, aiming to reduce the search space of possible combinations forming a pattern, to make it partially or fully retrievable.

Categories and Subject Descriptors

D.4.6 [**Software**]: Operating Systems—*Security and Protection*

General Terms

Security, Human Factors

Keywords

Android, smudge attacks, usability, pattern lock

1. INTRODUCTION

Nowadays, passwords are integrated in people's routines. Humans authenticate themselves using keyboards, fingerprint readers or touchscreens. Smartphones hold an important amount of information about the owner and for this reason people tend to lock them using the provided mechanisms. In most cases, phone lock mechanisms are implemented either as a PIN or a password.

Contemporary smartphones using the Android Operating System adopt a type of lock mechanism different to traditional PIN codes. This approach, called 'pattern lock', is based on existing research on graphical passwords [2] and requires the user to form a pattern on the screen by drawing lines in order to unlock the device. Its interface consists of 9 nodes in a 3x3 grid formation. Users start by touching one of the dots to make it the start point and swipe their fingers to add dots and form a pattern. However, there are some constraints while setting a scheme. It takes a minimum of 4 and a maximum of 9 dots to create one, each node can be visited only once and a previously not visited node becomes visited if it is on the way of a horizontal, vertical or diagonal line. Due to these constraints, the total number of possible patterns is 389,112 [1].

There are various types of attacks that can be used against a device to retrieve its pattern lock. Typical security attacks would entail attempts to exploit flaws in the theoretical design of a scheme or brute force a security mechanism. Brute forcing attacks against PINs or pattern locks may be rendered ineffective, if the number of unsuccessful attempts permitted is very limited and the device locks after that. Attacks that do not rely on brute forcing or exploiting a design weakness, but instead, are based on information gained from the physical implementation of a security scheme, are called *side channel attacks*. Some of such physical attacks against pattern locks aim to retrieve a pattern using physical traces left by the user, e.g. fingerprint marks left on the device's screen [11]. Others, such as *psychological attacks*, aim to detect user bias in PIN and pattern setting. Such information could be used to drastically limit the search space of possible combinations, in the same manner that heuristics about the use of meaningful passwords (e.g. familiar words) reduce the search space of a brute-force password attack [5].

In this paper, we attempt to combine physical attacks that relate to traces left from the use of a phone, with heuristics about the way users set lock patterns in order to facilitate attacks on this security mechanism. We use an optical camera and a microscope to analyse oily residues left on the screen, to evaluate the effectiveness of such relatively mid-term lived physical traces. We also exert a thermal camera to analyse heat traces left on the screen after drawing a pattern (shorter-term lived traces). To enhance the effec-

*Corresponding Author. Panagiotis Andriotis, Theo Tryfonas and George Oikonomou are with the Bristol Cryptography Group.

WiSec'13, April 17-19, 2013, Budapest, Hungary.

tiveness of the physical examination, we exploit outcomes of trends related to the setting of pattern locks. Therefore, we analyse the average pattern length, the number of direction changes when drawing patterns, the start points, end points and sub-patterns with length one to four nodes. To achieve this we conducted a pilot survey collecting data from 144 participants, indicating that there exist useful detectable trends when people try to form such a password. We therefore demonstrate that the combination of physical and psychological attacks can diminish the security efficiency of the pattern lock, revealing parts or the whole of a pattern.

The rest of this paper is organised as follows. In section 2, we discuss relevant research considering attacks on passwords. In section 3 we present our experiments to attack Android's graphical password scheme. In section 4 we present a preliminary evaluation of the proposed combination of physical and psychological-based attacks. The conclusion is drawn in section 5 and ideas for further work are also discussed.

2. BACKGROUND

2.1 Text-based and Graphical Passwords

Text-based passwords and PIN codes are normally coupled with bank accounts, computational devices etc. Individuals posses several accounts and numerous passwords that need to remember. Thus, the users often have to balance usability with security. As a consequence, they may recall another account's password [8] or even worse use the same across all their accounts [2]. If a word-like password is chosen, it may be possible for an attacker to retrieve it by using dictionary based attacks. On the other hand, if random characters have been set as password, it is highly likely that the user will fail to fully remember the sequence [8]. This renders text-based passwords hard for the legitimate user and easy for the attacker. Another aspect that makes the text-based passwords hard to remember is the way the human brain works. According to Dual Coding Theory, cognition is composed of two separate parts: nonverbal and verbal systems [4]. Having different systems in the brain to process the verbal and nonverbal information, humans perform differently in these two ways when it comes to remembering. Text requires an additional process of associating symbols with a contextual meaning [2].

Graphical passwords may come in much more variety compared to text-based solutions. They can include procedures such as clicking some points on an image, drawing a line or a shape. The most important advantage they provide is the possibility to define a password in a way that is memorisable by the user and yet, still hard to guess by the attacker. However, graphical passwords can also have their weaknesses, if we take into consideration the fact that users may select their graphical passwords with respect to some meaningful process. Thus, human psychology can be associated with the choice of a graphical password. Studies on image-based graphical passwords show that humans tend to choose popular points (called hotspots) on the image [10]. In their experiments, Thorpe and van Oorschot [10] argue that there are some general hotspots and areas on images that people tend to select. Furthermore, they are trying to answer the question if we can successfully build brute-force dictionary attacks on graphical passwords by defining weak password subspaces and applying attacks using complexity properties, such as password length, number of components, and symmetry [7]. Their predictive model leads to password rules and propose a set of precautions to increase security.

Another study provides 9 different face images to users and lets them choose 4 of them in a sequence to form a password [3]. Using this 'face selection' mechanism, they collect passwords from 79 participants. The results are significant and show that a number of passwords set by males can be easily guessed in 2 attempts. The fact that humans have similar preferences on graphical passwords provides reasonable grounds to investigate if there exist sub-patterns preferred by users when forming a pattern lock.

2.2 Methods of Pattern Lock Retrieval

Android's pattern lock mechanism relies on users swiping their finger to unlock the device. This action leaves behind an oily residue or smudges. Relevant research on retrieving lock patterns using standard optics is conducted by Aviv et al. [1]. In their work, they demonstrate how recovering smudges using a light source and a digital camera is possible due to the fact that touchscreen surfaces are reflective rather than diffusive. Experimenting with directional and omnidirectional light sources and testing angles ranging from 0° to 180°, by taking pictures at steps of 15°, they found out that the smudges were visible in most cases when a directional light source was used. Omnidirectional light sources prove to create a full reflection effect at all angles rendering this type of light source unusable. Apart from the ideal photograph capturing angles to retrieve smudges, the experiments focus on various states of a touchscreen such as: pattern entered using normal or light touches by the user, pattern entered before or after phone usage. Note that the notions of 'normal' and 'light' touches are not quantitative in this study, and thus must be intuitively guessed. For this reason, we assume the light touch stands for intentionally low pressure touches to minimise any smudge left behind, whilst the normal touch is the one made without any concern of leaving a smudge behind.

The smudge persistence of the patterns was tested on two phones. It is indicative that different touchscreen surfaces of Android phones (even from the same manufacturer) may have different properties with respect to capturing and retaining physical traces. When all angle setups are taken into consideration they derive that the best angle to retrieve a pattern is 60° with 80% of the lighting scenarios resulting in nearly perfect retrieval [1]. It is also noted that the directionality was discernible which is particularly important because it decreases the number of attempts to unlock the device. Their results highlight that intentionally cleaning with cloth or putting the phone to pocket was not enough to prevent pattern retrieval. It is important to mention that the researchers preferred describing the process as 'simple clothing', which probably means that the results may not hold true when the screen is rubbed thoroughly. Overall the optical method is particularly efficient as all it requires is a directional light source and a digital camera. An attacker can easily and quickly capture a photo of the touchscreen from a useful angle and perform any necessary contrast and brightness adjustments on the photo to retrieve one's pattern lock. As discussed in [1], even if the pattern is only partially retrievable, multiple photos taken in different times may reveal the full pattern.

The use of a thermal camera to retrieve the PIN codes

is an already existing attack on other devices such as ATM keypads (Mowery et al. [6]). Although there are two keypads for testing, one being metal and the other plastic, the tests are carried out on the plastic keypad as it is indicated that the metal keypad's conductivity renders it impervious to attack. Data gathered from 21 people and 27 PIN combinations display that the heat transferred to the keypad depends on the amount of pressure exerted to the keys as well as the warmth of the hand. However, the heat of the ATM was not taken into consideration as the keypad is used as an isolated test bed, without being wired to or placed on any electronic device. The thermal image shows a clear distinction between the background and the touched keys and right after the PIN entry, it displays with no hassle which buttons (and in which order) are pressed. An important aspect of thermal images is that they may be useful in situations where it is not possible to retrieve the smudges using standard optics, in a dimly lighted place, for instance.

3. EXPERIMENTS AND RESULTS

3.1 Physical Attacks

We tried to replicate Aviv et al.'s [1] methods using a different camera and smartphone. We used a Samsung Galaxy S featuring a Gorilla Glass screen (which is widely used by different manufacturers), a Panasonic Lumix DMC-TZ5 compact camera to conduct the attacks and a hard light source to achieve edged shadows. The objective in this section is to confirm previous work and for this reason the same person first draws the patterns and then conducts the attack in an open environment, with photos taken from a 60° angle. We performed an optical camera attack on three different surface conditions in terms of cleanliness. The first one was an attempt to retrieve the pattern drawn on a clean screen. It is evident that the pattern can be fully retrieved without any difficulty at this stage. The second test adds a light clean up to the first state. The 'lightness' of the clean up is indeed subjective and, in our case, we aim to mimic a person casually cleaning the device's screen without the specific intent of removing any oily residues. The oily residues turned to be quite resistant against simple cleaning attempts. Therefore, although the pattern can fade slightly, it can still be almost fully retrieved (some nodes might disappear). The final test was conducted on a heavily cleaned up surface. For this test we mimic a person determined to clean all the smudges on the screen at once. In this attempt most of the pattern is lost, except some diagonal lines. Due to increased entropy, it is also not possible to tell the directionality of the pattern. Therefore, we conclude that it is possible to capture patterns using compact optical cameras in most of the cases where the phone is not heavily used or cleaned and where there is efficient lighting.

For the microscope attack experiment, we used a USB microscope with 400x magnification. Our experiments followed the same logic we used for the optical attack. However, in the microscope case, we assume that the attackers have already seized the smartphone and are able to investigate the screen in a laboratory as long as desired. Similar to optical camera results, the microscope performed well during the first two cases. Lines and directionalities were very easily seen (full retrieval). There is, however, loss of some detail after the first clean up. For the heavy clean up case the microscope performed slightly better than the camera providing more details of smudge residues, but assuming that attackers can make use of a Digital Single-Lens Reflex (DSLR) camera in a controlled environment, it is highly likely to gather similar results without the need of a microscope.

The goal of the thermal image attack was to retrieve the pattern by examining the heat trace left by the finger on the device surface. The camera we used was a FLIR E30 and the experiments have been conducted from a distance of approximately 1 meter. The ambient temperature was 26°C, the light was low and no direct sunlight was coming to or near the device. Since time and heat are the main factors that contribute to form the results, the test cases were different than the previous: drawing a pattern on an idle device and drawing a pattern on a recently used device. The first scenario experiments revealed that it is possible to retrieve parts of a pattern via thermal imaging. We managed to observe the heat trace for 3 seconds after the pattern was drawn. However, we were unable to extract the pattern from a recently used device. When the device runs for a short period of time, its CPU starts to emit a considerable amount of heat. This in turn, heats up the upper and centre parts of the device rendering finger's heat untraceable. Even in the idle state, the CPU part of the device is considerably hotter. Consequently, the top three dots are hard to detect in most circumstances. To conclude, whilst it may not be a preferable attack compared to other options, a thermal attack might be used in the future, as the sizes of manufactured components diminishes and chip voltages are lowered.

3.2 User Tradeoffs between Security and Usability for their Choice of Pattern Locks

In order to study the effect of psychological or behavioural factors on pattern setting we conducted a web-based survey. This method was chosen because the participant does not need to own a specific smartphone. We used JavaScript, PHP, and AJAX to create the web-based survey and on the server side we held a MySQL database to store the given data. The pattern lock simulation utilised RaphaëlJS, which is a vector graphics library for drawing objects. The results presented in our work are calculated after filtering the database from irrelevant entries (144 unique participants). The survey started with basic demographics, continued with questions about participants' smartphone experience and their opinion on the notion of device locking and finalised after two pattern entries. The first was a pattern the user thought would be easy to remember and the second was a pattern that the user thought would be a secure password.

Summarising the findings of the survey we deduce that 65.97% were males and 34.03% females. The majority of the users were aged between 18 and 29 inclusive (81.25%) and the next more frequent age bracket was 30-49 (15.28%). This figure was expected because the survey was promoted through social media and through a university mailing list. 79.86% of the participants have owned a smartphone at least once, 92.17% of which still own a smartphone. Among them, 48.11% currently use iOS and 40.57% use Android. Symbian and Blackberry follow with 5.67% and 3.78% respectively. The people who ever owned an Android smartphone had at least one year of experience with it. 47.22% of the participants with a smartphone use any type of screen lock whereas 52.78% do not. The basic reasons they use a screen lock is to protect personal data and prevent others fiddling with the device. 65.98% of the participants believe that the

Table 1: Average pattern lengths and standard deviations.

Group	Average Length		Standard Deviation	
	Secure	Easy	Secure	Easy
Females	6.16	5.94	1.87	1.75
Males	6.89	6.32	1.91	1.94
Total	6.64	6.19	1.92	1.88

Table 2: Average number of direction changes (all users).

Average Changes		Standard Deviation	
Secure	Easy	Secure	Easy
3.57	2.74	1.65	1.59

highest risk that would compromise a lock is shoulder surfing. Smudges left on the screen and cameras in the room have the same rating of being the highest risk with 15.97%, yet the former has been selected more times as the second highest risk compared to the latter, rendering it the second highest risk after the shoulder surfing. Furthermore, 57.64% of the participants thought the secure pattern they entered is usable in everyday life, while 42.36% did not. Finally, 35.42% of the participants thought that the easy pattern they entered was secure enough, while 64.58% did not.

3.3 Secure Pattern Analysis

3.3.1 Pattern Length and Direction Changes

As part of our analysis, we calculated the average pattern length of the secure and easy patterns (and their standard deviations). The average length calculated by summing the number of dots used and dividing that value to the number of participants. While the average pattern length for a secure pattern drawn by a male participant is 6.88 dots, females averaged 6.16 dots. The same situation can be observed in easy patterns: the average among males is 6.32 dots while among females it is 5.94 dots. The total average lengths for secure and easy patterns are 6.64 and 6.19 respectively (Table 1). Results indicate that there may be a difference in perception of secure pattern length and direction between male and female participants. As part of our future work, we are planning to test the statistical significance of this claim, using a larger number of participants. Another indication of a pattern's security efficiency is the number of direction changes made per pattern (Table 2). We assume that for humans, a direction change is a more difficult move than following a direct line. Consequently, we deduce that when a user makes more direction changes in a pattern, then it gets more complex, hence more secure. The average number of direction changes made in a secure pattern is 3.57 and the average number of direction changes made in an easy pattern is 2.74. This finding demonstrates that secure patterns have more direction changes with respect to their lengths, rendering them more complex.

3.3.2 Entropy

In the following sections of survey data analysis, we calculated the Shannon's entropy while studying sub-patterns, start and end points for the secure patterns. For monograms, start and end points, entropy is calculated based on the probability of point X being selected in the pattern or being the start (or end) point. For N-grams, we calculated conditional entropy, whereby the probability of point X appearing N^{th} in the pattern is dependent on which N-1 points have been used in the pattern so far.

With that in mind, for conditional entropy calculations (F_N: N-gram entropy), we used Shannon's formula [9]:

$$F_N = -\sum_{i,j} p(b_i, j) \log_2 p(b_i, j) + \sum_i p(b_i) \log_2 p(b_i) \quad (1)$$

in which b_i is an (N-1)-gram (a pattern that consists of N-1 nodes), j is an arbitrary node (following b_1) that has not yet been chosen and $p(b_i, j)$ is the probability of the N-gram b_i, j. Note that in the case of sub-patterns, we consider $p(b_i, j)$ as $p(b_i||j)$, where $||$ stands for concatenation.[1] For instance, if the bigram is $b_i =$ "01" and $j =$ "2", then

$$p(b_i, j) = p(\text{"012"}) = \frac{\#\ of\ occurrences\ of\ trigram\ \text{"012"}}{sum\ of\ occurrences\ of\ all\ trigrams}$$

3.3.3 Start and End Points

An interesting observation from the survey is the way participants preferred to start their secure patterns (Figure 1a). More than half of them (52.08%) started their secure patterns from the top left node. The entropy of the start points is 2.35 bits compared to a maximum of $\log_2 9 = 3.17$ bits, for which all the dots must have the same probability. This imposes heavy bias and makes the first dot highly predictable. It is important to note that the survey did not examine whether the user is right-handed, left-handed or ambidextrous (we will examine that in subsequent iterations of the survey). Additionally, the survey could be filled either by using a mobile device or a computer, which means participants might have used a mouse to draw the pattern. Nevertheless, participants consistently chose the top left dot as the starting point. A reason for this clustering can be linked to participants' geographical positions. Most of the entries in our pilot run originate from across Europe and the United States. Most of these countries' native alphabets consist of Latin characters and consequently, their writing starts from the top left corner and ends in the bottom right. In addition, the survey's language is English, which may make the participants think in Latin language style even if they have a non-Latin based native alphabet. As a result, they may be inclined to start from the top left, because this looks like a more natural starting point. The collection of data from participants that have top-to-bottom or right-to-left native languages would provide interesting results in the future.

We then checked the ending dots for secure patterns. Even though there was no single dot on which most participants preferred to end their secure pattern, the bottom right was the most selected node with 20.83% (Figure 1b). The entropy calculated is 3.00 bits for the probabilities of end points. The ending dots were mostly focused on right and bottom. From this observation we deduce that the most frequent paths before the ending node can be found between the top right and the bottom right dot. As expected, these results also conform to the assumption about the Latin alphabet made on the analysis of start points.

3.3.4 Sub-patterns Analysis

One of the main objectives of the current work was to investigate the possibility to find reccurring sub-patterns

[1]Android's screen lock pattern nodes are represented with numbers starting with 0 from the top left node.

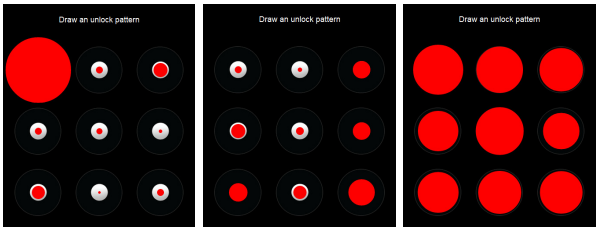

(a) Start points (b) End points (c) Monograms

Figure 1: Node usage (radius depicts frequency).

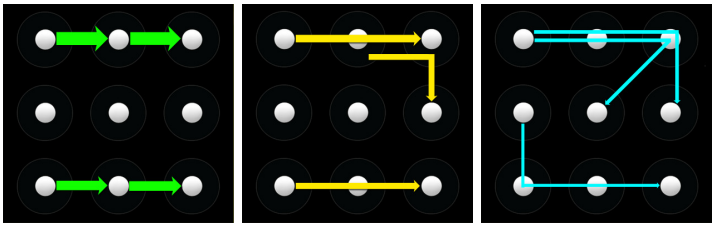

(a) Bigrams (b) Trigrams (c) Fourgrams

Figure 2: Most frequently drawn paths ('secure' patterns).

within the responses (focused on secure patterns). Extracting these sub-patterns would allow an attacker to guess a partially retrieved pattern's missing parts easier. In other words, an attacker can incorporate the physical attacks with the behavioural attacks to fully retrieve the pattern. The first step of our sub-pattern analysis involves monograms. We estimated the frequency of appearance of each dot to explore the existence of any particular nodes that are frequently chosen (Figure 1c). The result depicts that there is no significant bias towards any of the dots; they are more or less equally used in patterns, hence monogram entropy is 3.16 bits. We then looked for bigrams, a sub-pattern consisting of two dots. Since bigrams and other longer n-grams create a path, the directionality of the path is taken into account during analysis. There have been some bigrams that occurred especially frequently in patterns collected. In Fig. 2a, path width depicts the frequency of that particular bigram. In this case, the thickest path represents that 32.64% of the participants drew that bigram, while the thinnest represents 23.61%. Using Shannon's entropy, bigram entropy is calculated as 5.47 - 3.16 = 2.31 bits. The maximum entropy for the bigrams is $\log_2 72 = 6.17$bits. Out of 72 possible combinations 64, of them were drawn at least by one participant. Continuing with trigrams, the analysis shows that 18.75% of the participants drew a path from the top left dot to top right dot at one point of their patterns. The thinnest path in Figure 2b represents 14.58% of the participants. The trigram entropy is 6.99 - 5.47 = 1.32 bits. Maximum trigram entropy is $\log_2 504 = 8.98$bits. Out of 504 possible combinations, 203 were drawn at least by one participant. Finally, we conducted a four-gram analysis. Three four-grams stood out of the rest with two of them being drawn by 9.02% of the participants and the other being drawn by 7.64% (Figure 2c). Thus, it is easy to spot a trend towards left to right and top to bottom in these sub-patterns, which contributes to the assumptions made on the psychological behaviour the participants display. The four-gram entropy is 7.75 - 6.99 = 0.76 bits. Maximum fourgram entropy is $\log_2 3024 = 11.56$bits.

4. EVALUATION OF THE RESULTS

Our next step was to integrate our physical experiments and survey findings to propose a scheme, which could increase the success of such a combination of 'soft' and 'hard' side-channel attack on lock patterns. A common physical attack using an efficient optical camera combined with a psychological attack utilising the results of our survey should reduce the number of possible combinations and make pattern retrieval more efficient.

Table 3: Recovery of features.

Optical Attack	Number	Percentage
0 - 49% of pattern	5/22	22.73%
50 - 99% of pattern	5/22	22.73%
100% of pattern	12/22	54.54%
Total Recovery	18/22	81.82%
Phychological	**Number**	**Percentage**
Start point	18/22	81.82%
End point	11/22	50.00%
Bigrams	12/22	54.54%
Trigrams	7/22	31.81%
Fourgrams	4/22	18.18%
Direction (C)	14/22	63.63%
Total Retrieval	20/22	90.9%

To evaluate our proposed attack scheme we conducted a new experiment and derived data from a new set of 22 participants, which were not among those who took part on the web-survey. 15 of them (68.2%) were males and 7 (31.8%) females. 86.4% were aged between 21-30 and the age of the rest was 31-40 years old. The participants came from Europe (59.1%), Asia (31.8%) and America (9.1%). The experiment took place at a laboratory. They were asked to think of a secure pattern that they would probably use on their smartphones and then apply it on a real device. We copied and drew their patterns on paper and took photographs of the smartphone screen for further analysis. We marked the drawings with serial numbers before taking the photographs to ensure anonymity. The scenario we investigated assumes light usage of the phone after the pattern was entered, thus, before the photograph was taken we rubbed the screen gently on a cotton surface. We used an HTC Desire smartphone and a Nikon D40x DSLR camera for this experiment.

During the analysis of our data we investigated the correlation between the behavioural trends described in section 3 and smudges left on the screen. We used the following sequence. First, we set the reference standards for this experiment. The presented data in Figure 1 show that the 4 most preferred start points are the corners of the screen. In addition, the 4 most visited end points are those located at the right hand side and also the bottom left node. Finally, we took into account the most preferred N-grams (Figure 2). At the first step of the investigation the nodes and the edges of each pattern were recovered by scholastically analyzing the photograph of the given schema (optical). Then, we compared the findings of the former examination with our reference standards by looking at the drawing we had made

Table 4: Feature recovery of irretrievable patterns.

Physical attack	Number	Percentage
Start point	4/4	100%
End point	3/4	75%
Bigrams	3/4	75%
Trigrams	3/4	75%
Fourgrams	1/4	25%

for the specific pattern. The gathered information answered the question whether the pattern used a common start and end node and whether common N-grams have appeared. If at least one edge of the examined pattern was recovered either by the optical or the behavioral attack, then we can say that we achieved a partial retrieval. Table 3 demonstrates the results of our examinations.

The use of camera revealed, either fully or partially, 18/22 (81.81%) patterns. The psychological attack confirmed that 81.81% of the participants started their passwords using the most common start points and half of them ended their patterns at the expected end points. The average direction changes for males were 3.19, for females 2.83 and average pattern length for males was 7 and for females 6.33. We also saw some of the most common bigrams, trigrams and 4-grams presented in section 3 (popular bigrams were more frequent: 54.5%). Finally, the combination of the two attacks resulted in full or partial retrieval of 20/22 patterns and 100% of the patterns contained at least one of the reference standards.

This statistic shows that combining our web-survey results with well-known physical attacks we can increase the possibility to recover a pattern. We investigated the 4 patterns for which the optical attack did not reveal any information. Table 4 provides the results that justify the assumption that, even without any physical information, the psychological attack can still narrow-down the search space. One can argue about the sample size (4 patterns) but table 4 provides an indication that in most of the cases it is possible to retrieve parts of a pattern. At this table we present the success of the behavioural attack on the patterns that were not recoverable by optical attack. We can see that all of them contain at least one of the reference standards and specifically all contained an expected start point and also, most of them included at least one bigram, trigram and end point.

5. CONCLUSIONS AND FUTURE WORK

We successfully managed to attack an Android pattern lock using various physical attacks. We argue that currently an optical camera or a microscope are the best ways to perform physical attacks and produce quality results. Additionally, we have observed that humans tend to use specific heuristics when they form their lock patterns. We deduce that these heuristics are biased from aspects of peoples' context (e.g. spoken language). Our research demonstrated that it is possible to use the conclusions of our survey to increase the effectiveness of recovering patterns when combined with a successful physical attack. Further work has to be done to create a more global research, which will include other parameters that may be of significance, such as the user's educational level, geographical location and other demographic features. In our evaluation we underlined a trend to clockwise draw a pattern but this is an observation that must be further examined in the future. It would be also interesting to design a brute force attack model to allow a legitimate user to recover the pattern combining the artifacts found on the screen with the findings of the current research.

6. ACKNOWLEDGMENTS

This work has been supported by the European Union's Prevention of and Fight against Crime Programme "Illegal Use of Internet" - ISEC 2010 Action Grants, grant ref. HOME/2010/ISEC/AG/INT-002.

7. REFERENCES

[1] A. J. Aviv, K. Gibson, E. Mossop, M. Blaze, and J. M. Smith. Smudge attacks on smartphone touch screens. In *Proceedings of the 4th USENIX conference on Offensive technologies*, pages 1–7. USENIX Association, August 2010.

[2] R. Biddle, S. Chiasson, and P. C. Van Oorschot. Graphical passwords: Learning from the first twelve years. *ACM Computing Surveys*, 44(4):1–41, August 2012.

[3] D. Davis, F. Monrose, and M. Reiter. On user choice in graphical password schemes. In *USENIX Assosiation Proceedings of the 13th USENIX Security Symposium*, pages 151–163. USENIX Association, August 2004.

[4] D. J. Delprato. Mind and its evolution: A dual coding theoretical approach. *Psycological Record*, 59(2):295–300, September 2009.

[5] G. Fragkos and T. Tryfonas. A cognitive model for the forensic recovery of end-user passwords. In *Proc. of 2nd Intl. Workshop on Digital Forensics and Incident Analysis*, pages 48–54. IEEE CS Press, August 2007.

[6] K. Mowery, S. Meiklejohn, and S. Savage. Heat of the moment: characterizing the efficacy of thermal camera-based attacks. In *Proceedings of the 5th USENIX conference on Offensive technologies*, pages 6–6. USENIX Association, August 2011.

[7] P. C. v. Oorschot and J. Thorpe. On predictive models and user-drawn graphical passwords. *ACM Trans. Inf. Syst. Secur.*, 10(4):5:1–5:33, January 2008.

[8] M. A. Sasse, S. Brostoff, and D. Weirich. Transforming the 'weakest link' - a human/computer interaction approach to usable and effective security. *BT Technology Journal*, 19(3):122–131, July 2001.

[9] C. Shannon. Prediction and entropy of printed english. *Bell System Technical Journal*, 30(1):50–64, January 1951.

[10] J. Thorpe and P. C. van Oorschot. Human-seeded attacks and exploiting hot-spots in graphical passwords. In *USENIX Assosiation Proceedings of the 16th USENIX Security Symposium*, pages 103–118. USENIX Association, August 2007.

[11] Y. Zhang, P. Xia, J. Luo, Z. Ling, B. Liu, and X. Fu. Fingerprint attack against touch-enabled devices. In *Proceedings of the second ACM workshop on Security and privacy in smartphones and mobile devices*, pages 57–68. ACM, October 2012.