

(sp)iPhone: Decoding Vibrations From Nearby Keyboards Using Mobile Phone Accelerometers

Philip Marquardt et al.

ACM Computer and Communications Security 2011

n old kind of attack






How do we fix this problem?

Easy solution: users provide explicit permission



new kind of attack

We can use mobile phone accelerometers to detect vibrations




Related electrical/mechanical emanation attacks

Van Eck Phreaking (CRT electromagnetic emanations)


Tempest in a teapot: Compromising Reflections Revisited

Recreating key presses through a microphone




the new attack

Many users place their phones nearby when working on a computer
We can use this fact to our advantage to eavesdrop




the result of trying old techniques

We choose to use an iPhone 4 because it has a better accelerometer & gyroscope



...doesn't work



o what do we do next?

We choose to recognize key *pairs* instead of individual keys

We recognize whether keys are on the LEFT or RIGHT relative to a central line and if they are NEAR or FAR relative to some defined threshold distance α

For example, “Canoe” ->

LLN

LRF

RRF

RLF

Strings of length n can be split into $n-1$ abstract string representations

Creating our own neural network model

Step 1: Learning phase

Record each key press 150 times (total 3900 key-press events)

Create feature vector for each key drawing from x,y,z accelerations =>
<mean, kurtosis, variance, min, max, energy, rms, mfccs, ftt>

Word labeling: for each $n-1$ character pairs, concatenate random feature vectors for the corresponding keys

Can't be too specific -> to avoid overtraining, use even distribution of left, right, near and far labels



Creating our own neural network model

Step 2: Attack phase

Data Collection: Raw-acceleration data is collected

Feature Extraction: Feature-vectors are calculated

Key-press Classification: L/R labels and N/F labels are classified based on the neural networks

Word Matching: Words are matched against a dictionary and sorted; top scores are candidate predictions



How well does our model perform?

L/R classifier correctly identifies 91% of the time

N/F classifier correctly identifies 65% of the time

These percentages drop with more keypresses, which is to be expected



Experimental Results - Tests 1 and 2

Removed words of ≤ 3 characters

Test 1: 1 sentence -> 80% accuracy using first choice

Test 2: 10 sentences -> 46% using first choice, 73% within first two choices

1st Choice Correct = 80%
L/R Accuracy = 91.07%
N/F Accuracy = 70.15%

Typed Text: The birch canoe slid on the smooth planks

Recovered Text: *** punch canoe slid ** *** smooth planks

<- Test 1

1st or 2nd Choice Correct = 72.92%
L/R Accuracy = 83.95%
N/F Accuracy = 64.88%

Typed Text: Glue the sheet to the dark blue background

Recovered Text: Glue *** sheet ** *** well hogs background
blue

Test 2 ->

Typed Text: These days a chicken leg is a rare dish

Recovered Text: These days * chicken *** ** * rare dish

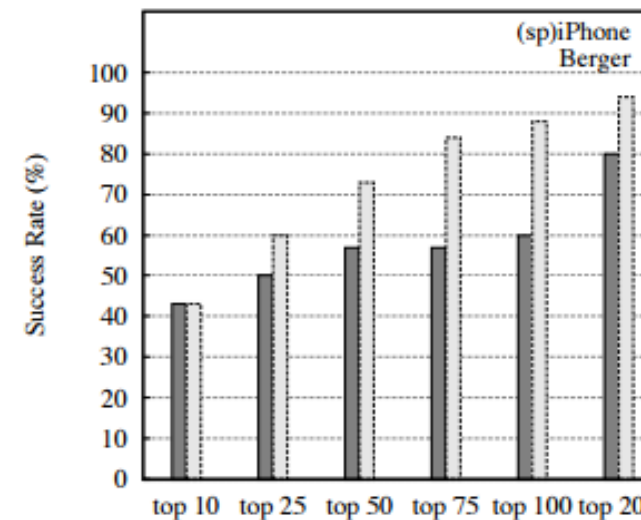
Experimental Results - Test 3

Comparison to previous work by Berger et al., using dictionary of 57,500 words and sentence with 4-9 characters per word

Berger: 43% accuracy within top 10 word guesses

Experimental result: 43% as well!

Experimental results less accurate than Berger when increasing the number of guesses...limitations?



Experimental Results - Test 4

A more realistic attack - USAToday article

Dictionary constructed using seven related news articles

40% in first choice, 53% in top 2, and 80% accuracy in top 5 predictions

1-5 Choice Correct = 80.00%

L/R Accuracy = 78.58%

N/F Accuracy = 61.09%

Typed Text: The Illinois Supreme Court has ruled that Rahm Emanuel is eligible to
run for mayor of Chicago and ordered him to stay on the ballot

Recovered Text: *** Illinois Supreme about *** ruled part wait Emanuel ** chicagos **
among Rahm eligible
might
night
Court
*** ** names ** Chicago *** printed *** ** look ** *** ballot
members
grinned
ordered

Challenges and Limitations

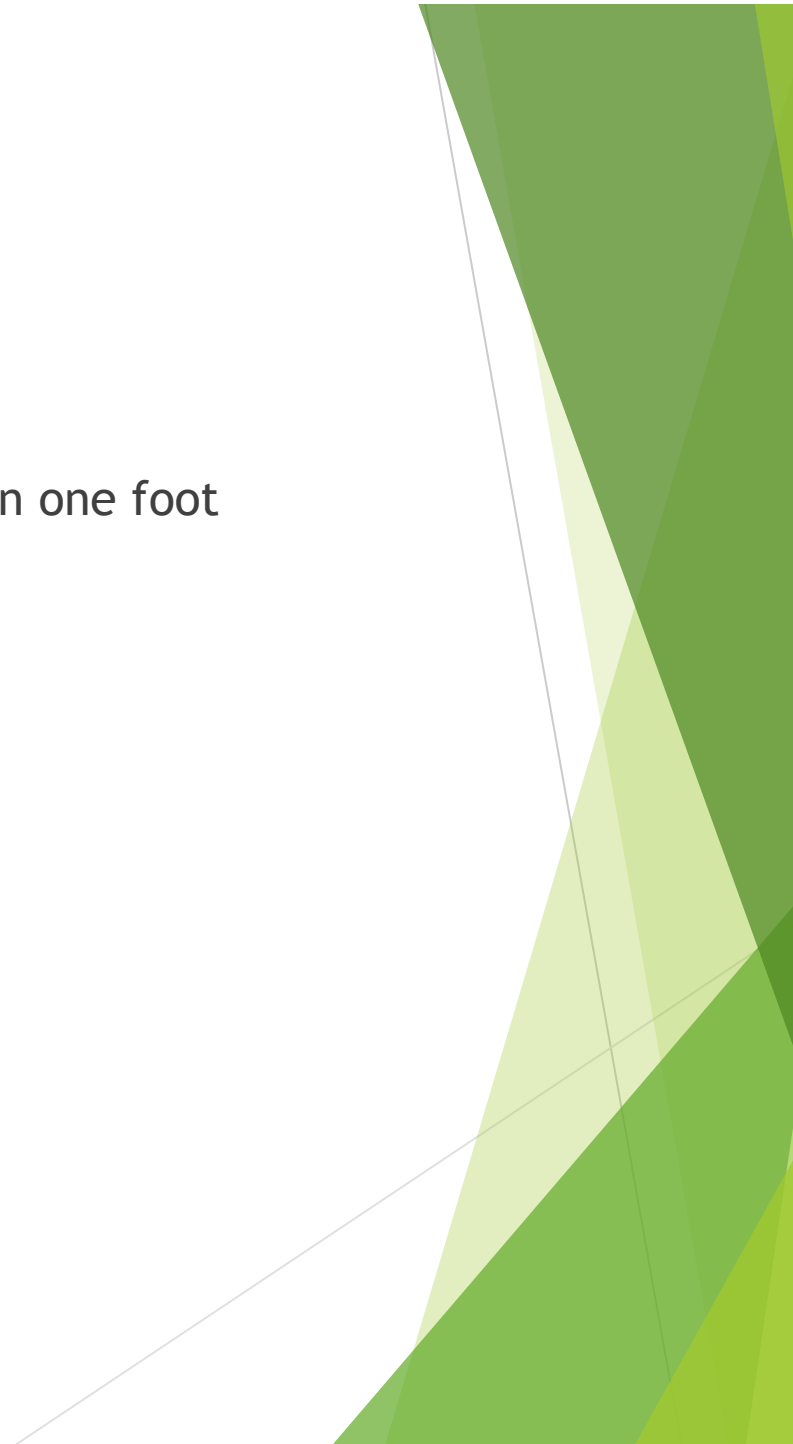
Distance and environmental factors - only sure to work within one foot

Orientation of the phone

Ambient vibration

Typing speed

Desk surface



How do we fix this vulnerability

Short term solutions

- Don't get too close!
- Permissions on accelerometer

Long term solutions

- Restricting data resolution to applications
- Being careful with all kinds of sensors in the future!



Discussion Points

Key contributions of the paper?

Limitations to this attack?

Is this paper relevant to other areas of security?

Thoughts on improving the accuracy/effectiveness of the attack?

What are ways we can combat these kinds of attacks?

