

```
|=====|
|=====|[ Sniffing Keystrokes With ]=====|
|=====|[ Lasers and Voltmeters ]=====|
|=====|
|=====|
|=====|[ By Andrea "lcars" Barisani ]=====|
|=====|[ <lcars_at_inversepath_dot_com> ]=====|
|=====|[ ]=====|
|=====|[ Daniele "danbia" Bianco ]=====|
|=====|[ <danbia_at_inversepath_dot_com> ]=====|
|=====|
```

--[Contents

- 0. DISCLAIMER
- 1. Introduction
- 2. Motivation
- 3. First Attack: Theory
- 4. The PS/2 Signal
- 5. Implementation
- 6. Data Analysis
- 7. Results
- 8. Attack Scenario and Workarounds

- 9. Second Attack: Theory
- 10. Implementation
- 11. Data Analysis
- 12. Results
- 13. Attack Scenario and Workarounds

- I. FAQ
- II. References
- III. Links

--[0. DISCLAIMER

All the equipment and/or circuits and/or schematics provided in the presentation must be treated as examples, use the presented information at your own risk, remember safety first.

--[1. Introduction

The exploitation of Electromagnetic Emanations and similar Side Channels has always been one of the most interesting and "exotic" areas of attack in the security field.

In the late 60's and early 70's the term TEMPEST[1] was coined to title an NSA operation which aimed to secure electronic equipment from leakage of compromising emanations. Well known TEMPEST research describes remote eavesdropping of CRT displays and most recently LCD displays, as well as optical emanations from appliances LED indicators.

Our research details two attacks, one against wired PS/2 keyboards, the other against laptop keyboards using respectively power line leakage and optical sampling of mechanical energy.

We describe how using relatively cheap homemade hardware we can implement basic but powerful techniques for remotely eavesdropping keystrokes.

--[2. Motivation

The two presented attacks partially builds upon existing concepts and techniques, but while some of the ideas might have been publicly hinted, no clear analysis and demonstration has ever been presented as far as we know.

Our goal is to show that information leaks in the most unexpected ways and can be indeed retrieved. If our small research was able to accomplish acceptable results in a brief development time (approximately a week of work) and with cheap hardware, consider what a dedicated team or government agency can accomplish with more expensive equipment and effort.

We think it is important to raise the awareness about these unconventional attacks and we hope to see more work on this topic in the future[2].

Last but not least.....hardware hacking is cool and everyone loves laser beams (this will make sense).

--[3. First Attack - Theory

The PS/2 cable of wired keyboards and mice carries the following wires:

```
-----  
- Pin 1: Data           / 6||5 \  
- Pin 3: Ground        | 4 || 3 |  
- Pin 4: +5 V DC       \ 2  1 /  
- Pin 5: Clock         -----  
- Pin 2/6: Unused
```

As the wires are very close and not shielded against each other it is theorized that a fortuitous leakage of information goes from the data wire to the ground wire and/or cable shielding due to electromagnetic coupling.

The ground wire as well as the cable shielding are routed to the main power adapter/cable ground which is then connected to the power socket and finally the electric grid.

This eventually leads to keystrokes leakage to the electric grid which can then be detected on the power plug itself, including nearby ones sharing the same electric line.

There might be other factors responsible in minor part for the signal interference like power fluctuations of the keyboard microcontroller, they are difficult to pinpoint but if present they can only augment the information leakage.

The clock frequency of the PS/2 signal is lower than any other component or signal emanated from the PC (everything else is typically above the MHz), this allows noise filtering and keystrokes signal extraction.

There has been some documentation suggesting the possibility of this attack in literature, though no extensive research is available. Recently a separate independent research which was developed simultaneously to our effort also suggests that "...the shared ground may acts as an antenna..." [3].

--[4. First Attack - The PS/2 Signal

The PS/2 signal represents an appealing and relatively favourable target for eavesdropping. The main advantage is its serial nature as data is transmitted one bit at a time, each keystroke is sent in a frame consisting of 11-12 bits (host-to-device).

As mentioned the clock frequency falls in the VLF (Very Low Frequency) category pulsing at 10 - 16.7 kHz range.

This is an example of what a PS/2 frame looks like:

```
-----  
|Start (1 bit)|Data (8 bits)|Parity (1 bit)|Stop (1 bit)|Ack (1 bit)|  
-----
```

The acknowledge bit is used for host-to-device communication only. As an example the letter 'b' (scan code 32) is the following frame:

```
---  
| 0 | 01001100 | 0 | 1 |  
---
```

--[5. First Attack - Implementation

In order to implement the attack the ground from a nearby power socket is routed to the ADC using a modified power cable (remember the disclaimer) which separates the ground wire for probing and includes a resistor between the two probe hooks. The current dispersed on the ground is measured using the voltage potential difference between the two ends of the resistor.

With "nearby" power socket we identify anything connected to the same electric system within a reasonable range, distances are discussed in the results paragraph.

In order to accomplish the measurement a "reference" ground is needed, as any ADC would need a proper ground for its own operation but at the same time the electrical grid ground is the target of our measurements. Because of this the main ground cannot be used as the equipment ground, as that would lead to null potential difference at the two ends of the probe.

A "reference" ground is any piece of metal with a direct physical connection to the Earth, a sink or toilet pipe is perfect for this purpose (while albeit not very classy) and easily reachable (especially if you are performing the attack from an hotel room).

We measured the potential difference on the ground wire routed from power plugs at 1, 5, 10, 15 meters from the actual target, due to the complex topology of the laboratory electrical system several junction boxes were present between the target computer plug and the sniffing device.

In all cases using a digital oscilloscope as an ADC, by sampling and storing the potential difference data, it is possible to obtain data about the ground wire "activity".

While the unfiltered signal apparently doesn't feature any useful information it was possible to successfully filter out the individual keystrokes from the original ground noise using the FIR filter. The PS/2 signal square wave is preserved with good quality (slightly modified by the anticipated artifacts introduced by the filter) and can be decoded back to the original keystroke information.

There has been no significant degradation of signal quality between the 1 meter distance test and the 15 meters one, suggesting that attenuation is not a concern at this range.

It should be noted that attenuation coefficients for wire copper are often estimated for much higher frequencies (>1Mhz) than the PS/2 signal, considering a typical copper cable with a coefficient of 0.1 dB after 60m theoretically (strong emphasis here) 50% of the signal survives. For reference a typical leakage emission has an output power of ~1 pW (10^{-12} Watt).

In conclusion the results clearly show that information about the keyboard keystrokes indeed leaks on the power grid and can be remotely detected.

We are confident that more expensive and sophisticated equipment can lead to much better measurements at a longer range.

--[8. First Attack - Attack Scenario and Workarounds

A good attack scenario for this kind of attack obviously involves the attacker being in a different room/area than the victim computer. In offices, houses, hotels it would be fairly easy to secure an attack spot with a power plug connected to the same electrical system as the victim room, possibly on the floor below or the adjacent room.

Other than diplomats, neighbours, ex-girlfriends and so on, it is worth to mention that an appealing category of targets are ATM/PoS machines and similar banking devices. Several ATM models in Europe are standard PCs with PS/2 (or similar) keypads and no strong electromagnetic leak shielding, depending on their location they are likely to share the same electrical system of the nearby shop or area.

The fact that the digits of the PIN code are the only input of the keypad narrows down the analysis required for retrieving it (of course we feel compelled to note that if the attacker has line of sight to the keypad it is more cost effective to simply point a zoom camera at the keypad).

The main workaround for this attack (other than obviously using laptops which are not connected to the power socket and have shielded power supplied anyway) is effective shielding of the RF emanations of the PC equipment. TEMPEST standards exist which define a series of protection requirements, but they haven't been completely declassified.

Extensive amount of tinfoil is not an effective workaround and it has been proved to make things worse in some scenarios.[4]

It is believed that USB keyboards are not affected by this attack as they use differential signaling for cancelling the noise, though USB microcontrollers within the keyboard are much more noisy than PS/2 ones and there is a chance that some fortuitous emanations might be present.

--[9. Second Attack - Theory

As the first attack does not work against laptops something different was needed for attacking this target.

Previous research[5] addressed using keyboard acoustic in order to mount a statistical attack for decoding the keystrokes, while these attacks are extremely fascinating we wanted to test something different that can be used at longer ranges.

Laser microphones are well known monitoring devices that can detect sound at great distances by measuring the mechanical vibration of glass windows (which resonate due to the sound waves that hit them).

The theory is that the mechanical vibration produced by keystrokes propagates on the laptop case carrying information that can be used to decode them. A laser microphone can be pointed at the laptop case directly instead of a window in order to sample those vibrations in a fashion similar to sound detection (effectively making the laser microphone a laser "vibrational sampler" as no sound is involved).

--[10. Second Attack - Implementation

While several commercial laser microphones are available at a high price, it is fairly easy to build your own for as little as 80 USD.

Here's the basic needed equipment:

- 1 x Laser
- 1 x Photoresistor or Photodiode
- 1 x Variable resistor
- 1 x AA Battery
- 1 x Universal Power Adapter
- 1 x Jack Cable
- 1 x Laptop with sound card
- 2 x Tripod
- 1 x Focusing lens (for long distances)

Optional components can include an amplifier and/or an optical bandpass filter.

We built a basic laser device with a cheap Class IIIR laser (670 nm, <5 mW power, <2 mrad convergence) slightly better than the average laser pointer. A photodiode works better than a photoresistor because of its increased response speed, example photodiode models are BPW21R and BP103.

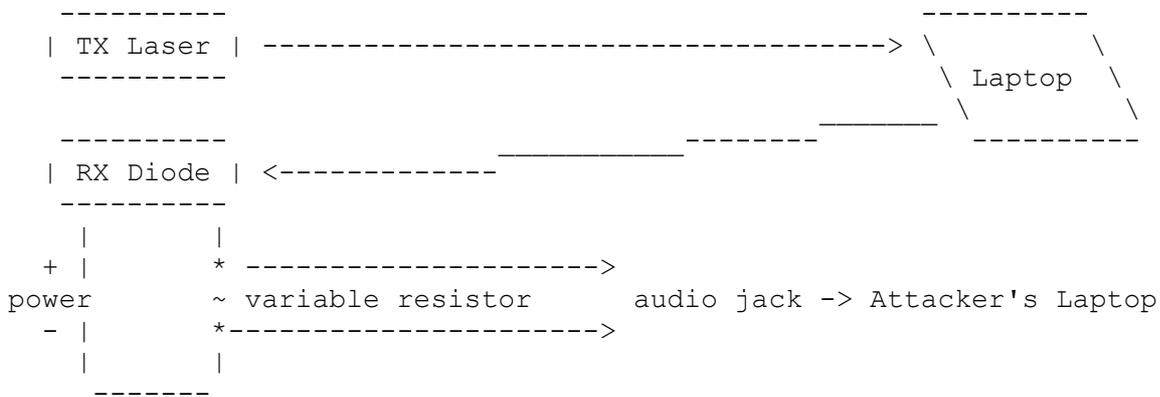
The laser device components are a transmitting side (TX), consisting of the laser, and a receiving end (RX), consisting of the photoresistor/photodiode which is routed to a standard laptop pc sound card using an audio jack cable.

The vibration of the target (in this case a laptop) results in movements of the reflected laser beam, this modulation is converted by the receiving side in an electrical signal which can be turned to digital data using the laptop audio card (which acts as a low cost ADC).

The power output level of the signal, using an AA battery, is compatible with standard sound cards but we recommend to test it with a voltmeter before connection (again, remember the disclaimer). In general it is pretty easy to saturate the device as the sensitivity of the receiving side is very high, this is why the resistor is needed in order to tune the circuit output power. The laser reflection is generally so powerful that the outer area of its circle is sufficient enough for the measure, while the bright center saturates the circuit.

The high intensity of the laser allows the receiver to distinguish the laser during day light and/or at longer ranges, while at night time the measurements are even more precise because no background light noise is present.

Diagram:



In order to test if the assembled device works correctly a good method is using it as a "normal" laser microphone against a window, if the device is tuned for detecting audio it will be good enough for vibration pattern detection.

--[11. Second Attack - Data Analysis

The vibration patterns received by the device clearly show the separate keystrokes, this means that previous research that involves analyzing the timing of the keystrokes can be reused with this method.

In addition, as the vibrational information is precise enough, we can compare the patterns to each other in order to assess the likelihood of the different keystrokes being the same (or different). This allows recovery of recurrent/distinct letters within the words and eventually the entire text which is being typed.

As the space bar is a key shaped in a substantially different manner than any other key on they keyboard layout, it is immediately possible to separate the words from each other. This greatly helps the data analysis as by making assumptions on the language which is being typed it is possible to narrow down the odds of small words and re-use that information throughout the analysis (as an example 3 letter words in English are likely to be either 'are' or 'the').

As the different vibration patterns are not going to be identical because of difference in typing speed and mechanical propagation a scoring technique is necessary for the comparison. Dynamic Time Warping (DTW) is a good old-fashioned technique for measuring the similarity of signals with different time/speed, it is generally applied to audio and video but in principle can be used with any signal.

More modern statistical techniques exist, like Hidden Markov Model (HMM), and can be surely employed with the same, if better, effectiveness.

It is important to emphasize that this attack doesn't requires previous knowledge or training about the victim (other than the language) as we perform the statistical comparison between the different keys of the same sniffed data. Knowing the context of the text it is possible to considerably narrow down the options with just a few words of data.

Additionally the order of the typed sequences is not a factor, as an example if someone types a password and then a page of text the latter analysis can be used to narrow down the options for guessing the password.

--[12. Second Attack - Results

From a signal detection point of view it was possible to obtain good results below 30 meters without any heavy tuning, using the cheap laser. Longer distances requires precise calibration and filtering and of course the more money is thrown at the laser quality the better the range is going to be.

Aiming the beam directly at the laptop case, generally the LCD display lid, proves to be effective. The top of the lid catches more resonant vibrations (to be subtracted later via signal analysis) while aiming closer to the hinges produces better results.

Here's a sample result dump from a pessimistic case scenario of just two words being typed:

chars 1 <> 7 = 0.066*	chars 7 <> 8 = 0.029*	chars 8 <> 7 = 0.029*
chars 1 <> 8 = 0.072*	chars 7 <> 1 = 0.066*	chars 8 <> 1 = 0.072*
chars 1 <> 3 = 0.167	chars 7 <> 3 = 0.161	chars 8 <> 3 = 0.146
chars 1 <> 10 = 0.188	chars 7 <> 10 = 0.191	chars 8 <> 6 = 0.226
chars 1 <> 6 = 0.209	chars 7 <> 6 = 0.270	chars 8 <> 10 = 0.244
chars 6 <> 10 = 0.160*	chars 10 <> 6 = 0.160*	chars 11 <> 1 = 0.065*
chars 6 <> 1 = 0.209	chars 10 <> 7 = 0.191	chars 11 <> 8 = 0.029*
chars 6 <> 8 = 0.226	chars 10 <> 1 = 0.188	chars 11 <> 7 = 0.072*
chars 6 <> 7 = 0.270	chars 10 <> 8 = 0.244	chars 11 <> 3 = 0.146
chars 6 <> 3 = 0.343	chars 10 <> 3 = 0.250	chars 11 <> 6 = 0.226

The lower the score the better the match. Characters 1, 7, 8 and 11 are definitely identical like 6 and 10 while characters 3 and 4 looks different than anything else.

Knowing where the space bar was we can group the different keys with the following pattern of 1?XY321 1321.

Here's what happens if we input the result to a very simple application that performs regular expression pattern matching against a dictionary using the supplied grouping.

```
$ ./WoF '1_XY321 1321' /usr/share/dict/american-english
```

```
hogwash hash  
salmons sons  
secrets sets  
sermons sons  
sockets sets  
soviets sets  
statues sues  
straits sits  
subways says  
tempest test  
tidiest test  
tiniest test  
trident tent
```

We can see that knowing the context it is immediately possible to assess that 'tempest test' and maybe 'secrets sets' are the most probable answers, and indeed the former is the correct one.

Adding an article to the phrase (like 'the') narrows down the options to just two possibilities. With a full page of text, while the matching process takes more time, it is easily possible to recover the entire text.

--[13. Second Attack - Attack Scenario and Workarounds

Obviously a line of sight is needed, either in front or above the target, for mounting the attack. While this is not trivial to achieve it is reasonably possible if the target is facing a window on a high floor or placed on a table in a location (an outdoor area as an example) where the attacker can reach higher grounds. The transmitting and receiving sides can be at two completely different locations.

A reflective area is needed for the attack and we found out that almost every laptop has a usable area. In case of IBM Thinkpads the logo on the lid can be used as well as the reflective plastic antenna for later models, Asus netbooks lid is entirely reflective and hence perfect for the attack. Apple laptops can be targeted on the Apple logo itself or, if you are attacking from behind, on the ultra-glossy screen.

Additionally it is possible to aim the laser at any reflective object present on the laptop support like glasses close to the laptop and so on, if the table is sufficiently elastic to propagate the vibrations the attack is successful.

While one laser device was used in our tests it is possible to combine more of them and have 2 or 4 devices aiming the same laptop simultaneously, it is also possible to use different kind of laser microphones that use interferometry in order to assess the Doppler effect of the frequency shift caused by the vibration. All of this can greatly help the measurement for longer ranges.

Stealthiness (as red laser dots on your laptop case might look suspicious now) can be easily achieved by using an infrared laser/receiving diode, though it might require an infrared camera or temporary guidance with a visible laser for the actual targeting.

The attack is possible even with a (possibly double) glass window in the way as reflection loss is ~4% at every pass.

As a workaround (other than avoiding the line of sight with the attacker in the first place) the only ways we can think of are using an extremely firm laptop (we have yet to find a model which satisfy this requirement), radically change position while typing every second or so (you might look weird while doing this) or "pollute" the data with random keys somehow and delete them with backspace afterwards.

--[I. FAQ

1. Where are the pretty pictures? I see only ASCII art here.

Check the links section down below for the full pdf presentation with all the pictures.

2. In the first attack can you detect different keyboards being used on the same electric line?

Yes, the PS/2 frequency is a range and it is very difficult to find two keyboards at exactly the same frequency. Unless you have thousands of keyboards it will be possible to differentiate them.

3. In the second attack does the result change if different people are typing?

Yes it does, in the sense that every person typing style will produce different vibrational patterns even for the same laptop. At the end though this is not a factor for the attack success as the analysis is assumed to be performed for a data set coming from the same person.

It is not possible to re-use the scoring from one person against a different one (unless we are talking about two identical evil twins)

--[II. References

- [1] - TEMPEST is believed not to be an acronym though several non-official ones have been proposed, the most catchy are "Transmitted Electro-Magnetic Pulse / Energy Standars & Testing" and "Tiny Electro-Magnetic Particles Emitting Secret Things."
- [2] - While drafting this whitepaper news broke out about a very interesting research on this same field, be sure to check out "How Printers can Breach our Privacy: Acoustic Side-Channel Attacks on Printers"
<http://www.infsec.cs.uni-sb.de/projects/printer-acoustic>
- [3] - Martin Vuagnoux, Sylvain Pasini (awaiting peer review at July 09)
"Compromising radiation emanations of wired keyboards"
- [4] - Ali Rahimi, Ben Recht, Jason Taylor, Noah Vawter "On the Effectiveness of Aluminium Foil Helmets: An Empirical Study"
<http://people.csail.mit.edu/rahimi/helmet>
- [5] - Dmitri Asonov, Rakesh Agrawal (2004) "Keyboard Acoustic Emanations"
Li Zhuang, Feng Zhou, J.D. Tygar (2005) "Keyboard Acoustic Emanation Revisited"

--[III. Links

- Project directory
<http://dev.inversepath.com/download/tempest>

|=[EOF]=-----=|