ON THE IMPLICATIONS, THE IDENTIFICATION AND THE MITIGATION OF COVERT PHYSICAL CHANNELS

Michael Hanspach[1] and Jörg Keller[2]

[1] *michael.hanspach@fkie.fraunhofer.de*
Fraunhofer FKIE, Wachtberg (Germany)

[2] *joerg.keller@fernuni-hagen.de*
FernUniversität in Hagen (Germany)

Abstract

Covert physical channels use physical means like optical emissions or acoustic wave propagation to connect isolated operating system compartments within the same computing system and independent devices that are physically separated by air gaps. We extensively discuss the implications, the identification and the mitigation of these covert physical channels. For the purpose of identifying covert physical channels during the design and evaluation of the computing system, an adaption of Kemmerer's shared resource matrix, namely the *physical environment matrix*, is presented. The physical environment matrix enables the computing system's designers and evaluators to systematically describe and classify physical environments and the associated covert channels that might be possible between two specific devices or operating system compartments. Through the specification of limited access to physical environments, the presence or absence of a covert physical channel can be formally demonstrated during the design and evaluation of the computing system.

Keywords: Malware, operating system security, covert channels, ultrasonic communication.

1 INTRODUCTION

Even in a security hardened computing system that might be running on top of a formally verified micro kernel running on top of a formally verified hardware, and utilizing a well-written security policy, we might be unable to preclude information leaks over covert physical channels, unless these covert channels can be systematically identified and mitigated. These covert physical channels could, for instance, be based upon light emission or on acoustic wave propagation, and even molecular communication [1] would be possible with potential multi-sensory computing systems (e.g., future smartphones).

While the term *covert* in *covert channel* reflects the matter of fact that the computing system's designer did not actively construct the covert channel [2], the covert channel's degree of *stealthiness* is used in this article to describe the difficulty of detecting the covert channel. Stealthiness is far from being an absolute property, as a covert channel that might be hard to detect for a human observer might sometimes be easily detectable for a specially prepared computing system.

Recently, covert physical channels have been discussed in the context of covert networks that consist of malware-infected computers and enable attackers to bridge air gaps [3]. In addition to this capability, these types of covert channels are equally suggested to be capable of breaking domain separation in secure operating systems [4]. In an operating system context, a covert physical channel uses the computing system's environment to transmit physical signals from one compartment to another using resources for physical communication that have not been designed for communication in the first place.

Computing systems within high-security environments operate in a narrow field between sharing and isolating resources. While we want to pool our computing resources and reduce the number of devices used for processing different types of sensitive data (e.g., in multilevel security), we have to be very careful with granting access to shared resources and the shared environment, because any shared medium could possibly be used to establish a covert channel that breaks the system's isolation policy. In this respect, the computer's physical environment deserves attention as it is necessarily shared by all operating system compartments and nearby computing systems. Until now, the computer's physical environment has not yet been systematically treated as a resource to be considered for isolation in computing system design.

Our contribution to the field consists of a systematic treatment of the implications, the identification and the mitigation of covert physical channels. We extensively discuss the threat model of different types of covert physical channels. Moreover, we present the *physical environment matrix* that enables designers and evaluators of computing systems to systematically identify and, by this, to mitigate these covert physical channels.

The remainder of this article is structured as follows: In Section 2, we give an overview about related work and how it differs from our work. In Section 3, we discuss the implications and the threat models of different types of covert physical channels. In Section 4, we demonstrate the systematic identification of covert physical channels between devices and operating system compartments. In Section 5, we briefly describe mitigation strategies against covert physical channels. In Section 6, we conclude the article.

2 RELATED WORK

Side channels over electromagnetic emanations (van Eck phreaking [5]) have been widely discussed in the past. Different types of emanations have been discussed in this context. Information leakage from optical emanations is studied by Loughry and Umphress [6], and acoustical side channels are discussed by Genkin, Shamir and Tromer [7], and by LeMay and Tan [8]. In contrast to side channels, covert physical channels are not the result of an unintended side effect in computing, but they are actively initiated by attackers through the misuse of environmental sensory devices.

Hasan et al. [9] present different types of physical emanations for use in command-and-control messages of botnets. In a different context, Frankland [10] describes an attack where keyboard LEDs are used to transmit a message, which is then recorded by a camera. In contrast to these authors, we develop a general theory of covert physical channels, where we target both covert physical channels for computer-to-computer communication and operating system covert physical channels for communication between separated compartments of the same computing system.

Kemmerer introduces the shared resource matrix that is designed to identify covert storage and timing channels within operating systems [11]. The shared resource matrix approach is very successful in identifying covert channels over shared resources. However, as covert physical channels are not established over *shared* resources, but over *cooperating* resources (see Section 4) that interact within a shared physical environment, a different approach is needed for covert physical channels. The *physical environment matrix* (as presented in Section 4) aims specifically at the identification of covert physical channels. Both local operating system covert channels (on the same device) and remote covert channels (between different devices) are targeted by the physical environment matrix.

3 IMPLICATIONS OF COVERT PHYSICAL CHANNELS

3.1 Infection methods

The attacker needs to infect a computing system over one of different ways in order to exfiltrate data: The attacker might plant a backdoor within the target's computing system in order to enable the covert capabilities right from the start. Alternatively, he might use removable media to spread the infection. Removable media have been shown to bridge air gaps by successively connecting them to different separated computers (as demonstrated with Stuxnet [12]). The attacker might strategically place an infected removable medium in the physical or social environment of the computing system, or he might infect the (previously trustworthy) removable medium over the Internet when it is connected to a vulnerable Internet-connected computer.

3.2 Threat model: covert networks

In a high-security context, computing systems are regularly separated by air gaps. Covert networks based on covert physical channels enable attackers to bridge these air gaps between infected computers. Gathered data might be further spread when a user carries infected mobile devices into different environments. These covert networks might also be connected to the Internet if one of the nodes offers Internet access as a service to the infected nodes. The attacker's computer needs to be temporarily in physical communication range to one of the nodes, or the covert network needs to be connected to the Internet to enable data exfiltration or remote manipulation.

Attackers might utilize covert networks in the following scenarios:

1. In a *targeted attack*, an attacker desires access to a specific computing system that is physically separated from other computers by air gaps. The covert network would be operating in a stealthy state, i.e., only occasionally forwarding data to other computers in order to avoid early detection.
2. In a *botnet-related scenario*, an attacker desires access to as many physically separated computing systems as possible, but he does not necessarily target a specific system (i.e., he does not know what information he might find where). The covert network would regularly forward data to other computers in order to maximize the impact of the network (leading to a less stealthy state).

3.3 Threat model: operating system covert channels

Traditional types of operating system covert channels have been extensively discussed [13]. Covert physical channels can be fully utilized as a replacement method for the prevailing applications of these types of covert channels.

An attacker with access to compartment p_j (this terminology is also explained in [4]) can utilize an operating system covert channel to exfiltrate data from a separated compartment p_i on the same computer. The attacker might gain control to compartment p_i over the previously described infection methods (e.g., removable media). After the takeover of the compartment p_i, the attacker can exfiltrate data from p_i at any time. If p_j is connected to any network (e.g., the Internet), data can first be exfiltrated over the covert physical channel and then forwarded to the network. Additionally, an insider attack can be conducted, where the attacker is granted access to the compartments, but lacks a method of exfiltrating the data from one compartment. This threat is especially relevant in high-security computing systems, where other types of covert channels might already be mitigated (see also [14]).

However, the threat could also be applied to other security and privacy scenarios, for instance, to connect isolated mobile phone apps. While it might not always be possible to gain control over both compartments p_i and p_j, the operating system covert channel

approach offers the advantage of a small and fixed distance between the sending and receiving device in comparison to the covert network approach. Because of this, the movements of the computing device are irrelevant to the availability of the covert channel, and the devices are always in communication range. Furthermore, the smaller distance between the devices can lead to a much higher channel performance as less communication redundancies and error correction measures are necessary.

The covert physical channel might be implemented as a unidirectional or as a bidirectional channel. In case of a bidirectional covert channel, bit errors could be treated by a higher-layer protocol supporting acknowledgment and retransmission of frames. In case of a unidirectional covert channel; bit errors could be treated with a frame structure containing the file ID, the file pointer (denoting the byte position in the file) and a CRC checksum (for detecting bit errors at the receiver side) in addition to periodical retransmissions of the file contents.

Because of this, unidirectional implementations of a covert physical channel would be most fitting to periodical transfer of small-sized text files containing critical information (such as private encryption keys associated with one compartment) and to transfer of larger files containing fault-tolerant data (e.g., some graphics formats). So, in respect to the fields of application, the unidirectional covert physical channel is similar to side channel attacks. However, in difference to side channel attacks, the covert physical channel is actively triggered by malware and not a result of unintended side effects in the computing system.

More advanced erasure codes like LT codes [15] might be utilized to further increase the reliability of unidirectional covert physical channels. In contrast to unidirectional implementations, bidirectional implementations of a covert physical channel would also be appropriate for transfer of larger files (e.g., containing classified data).

3.4 Implications on multilevel security

In an attempt to break the mandatory access control policy of multilevel security operating systems, an operating system covert physical channel could be used to connect a SECRET classified and a CONFIDENTIAL classified compartment, facilitating a data leak from SECRET to CONFIDENTIAL that would be prohibited by the implemented access control policy under other conditions. Even in system setups with equally classified compartments (e.g., SECRET and SECRET), directly connecting two compartments in a pure isolation hypervisor (generally disallowing communication between compartments, see also [16]) would lead to a violation of the mandatory access control policy. The reason for implementing this type of policy can be explained by the desire for need-to-know separation (i.e., access to classified data shall be granted separably, even among equal classification levels).

3.5 Implications on undesired communication between users

If we are looking at communication between users that are associated with computers or compartments, the precondition of malware infection dissolves. We consider a scenario with two insiders that control their own computing resources and that are not allowed to connect their computing resources. While both users lack established communication methods, they might be able to communicate via a covert physical channel.

4 IDENTIFICATION OF COVERT PHYSICAL CHANNELS

Countermeasures against covert channels generally include the limitation of resource sharing. The shared resource matrix has been successfully used for covert channel analyses of operating systems in the past [11]. The shared resource matrix is a

methodology for the identification of shared resources, which could possibly be used for establishing covert channels between the sharing compartments. For the identification of a covert physical channel between remote or local compartments, however, the shared resource matrix is not directly applicable, as it aims at shared resources, while we are looking at covert physical channels that build upon a shared physical environment with *cooperating* (but not necessarily *shared*) resources.

Cooperating resources are resources that are able to input or output a type of physical emanation (e.g., optical or acoustical emanations). Only resources that have not yet been established as a communication device are targeted by our analysis. Therefore, communication over established types of network interfaces (e.g., WLAN or Bluetooth) is not considered by this analysis. We present the physical environment matrix (Tab. 1), which was created with the shared resource matrix's philosophy of systematically identifying possible covert channel means in mind.

p_j \ p_i		Audible Sound			Ultrasound			Light	
		MIC	SPK	SPK-F	MIC	SPK	SPK-F	CAM	LED
Audible Sound	MIC	C	C	C					
	SPK	C	C	C					
	SPK-F	C	C	C					
Ultrasound	MIC				C	C			
	SPK				C	C			
	SPK-F								
Light	CAM								C
	LED							C	

Table 1: The physical environment matrix is introduced as a systematic methodology to analyze covert physical channels between isolated compartments.

The horizontal axis of the physical environment matrix features the accessible physical environments and the correspondingly accessing devices of compartment p_i, while the vertical axis does the same for compartment p_j. Both p_i and p_j are defined as distinct compartments (e.g., virtual machines or isolated application containers) on a computing system or on different computing systems. Cooperating devices are marked with a *C* in any inner cell of the matrix, where both cooperating devices are actually able to communicate unidirectionally or bidirectionally.

If a device is inaccessible by any compartment, the corresponding row or column could be removed as there would be no devices marked as cooperating devices in the appendant cells. To demonstrate the different capabilities of physical environment in this example, we split up the physical environment *sound* into *audible sound* and *ultrasound*. This allows us to represent devices that differ in their behavior regarding audible sound and ultrasound such as a filtered sound device. To show another similar application of the physical environment matrix, we introduce *light* as a physical

environment, which can be accessed by a camera or a LED, as light has been shown to be efficient in covert communications [10].

Regarding the types of devices, *MIC* is defined as a microphone, *SPK* is defined as a speaker, *CAM* is defined as a camera and *LED* is simply an LED. Finally, *SPK-F* is defined as a speaker with a filter that lets only audible sound frequencies pass through. All accessible devices (this might, for instance, also include motion sensors and temperature sensors) have to be considered in the physical environment matrix. This also includes indirectly accessed device capabilities, e.g., the capability to manipulate keyboard LEDs.

As it is shown in the physical environment matrix, *MIC* and *SPK* are always able to cooperate in the audible sound frequency range. *MIC* and *SPK* should also be able to speak to another distinct instance of *MIC* or *SPK*, respectively. This is because microphones and speakers are commonly known to be designed by very similar working principles, and a speaker might be wired to also act as a microphone and vice versa. If a compartment p_i claims control over both audio input and output devices (e.g., in IP phone setups), p_j might be able to access these devices as well by the means of I/O virtualization (e.g., as supported by VirtualBox).

In the shared physical environment *light*, *CAM* and *LED* are able to cooperate, using a unidirectional information flow in the form of *CAM* → *LED* (from left to right). Finally, there is an important difference in the ultrasound frequency range, as *SPK-F* is not able to cooperate with any other device, as the filtering effectively prevents ultrasonic communications. Physical environment matrices are formally defined by Eq. (1) where *P* is the set of compartments, *D* is the set of devices, and *E* is the set of environments.

$$\begin{aligned} &\forall p \in P, \forall d(p) \in D : \exists e(d(p)) \in E \\ &\forall p_i, p_j \in P (i \neq j) : \\ &\qquad \forall e(d(p_i)), e(d(p_j)) : \exists c \in \{C, \neg C\} \end{aligned} \quad (1)$$

Thus, every compartment *p* is associated with its own set of devices that are able to output a physical emanation. Each of these devices *d(p)* is associated with a physical environment *e(d(p))*. For each pair of these environment/device associations from two different compartments, the devices might be able to cooperate or not to cooperate. As shown, this procedure might be performed for every pair of compartments p_i/p_j from the available compartments *P*.

To enrich the semantics of the physical environment matrix, each cell might also contain the maximum performance of a potential covert physical channel. The maximum performance for the cooperating capabilities $C(p_i, p_j)$ of all potential covert physical channels between two compartments p_i and p_j is given by Eq. (2), where $C(p_i,p_j,k)$ is the maximum performance of the sending or receiving device in row or column *k* of the matrix.

$$C(p_i, p_j) = \sum_{r=1}^{R} \sum_{l=1}^{L} min(C(p_i, p_j, r), C(p_i, p_j, l)) \quad (2)$$

$C(p_i, p_j)$ is calculated by summing over all rows *1...R* and columns *1...L* of the physical environment matrix for p_i and p_j. In each row and column combination, the maximum performance of a potential covert physical channel equals the minimum of sending and receiving performance.

The physical environment matrix is specifically designed to be used for the design and evaluation of a computing system. During all parts of the product life cycle (i.e., before, during and after systems development), the physical environment matrix might be applied to identify the hidden communication capabilities of a computing system. The

physical environment matrix is not specifically designed to identify capabilities that might be added to a specific computing system during operation in the field (e.g., malicious hardware implants). Therefore, additional measures might be implemented to ensure the integrity of the secured computing system, and the physical environment matrix can only precisely describe the hidden communication capabilities of an integrity state of the computing system.

5 MITIGATION OF COVERT PHYSICAL CHANNELS

5.1 Mitigation strategies

From looking at the physical environment matrix, covert physical channels might be mitigated in the following ways:

1. Access to a physical environment might be completely prevented.
2. Access to a physical environment might be limited to humanly detectable signals.
3. Access to a physical environment might be limited to non-steganographic signals that are easily detectable with computer-aided detection methods.
4. If the presence of potential covert physical channels is accepted, further means might be implemented to detect communication over the covert physical channels.

To implement the approaches 2) and 3) within the physical environment matrix, each physical environment might be split up into humanly detectable and non-detectable environments (e.g., as demonstrated for audible sound and ultrasound in Tab. 1), and into steganographic and non-steganographic environments. Regarding 3), steganographic communication might be implemented by attackers that do not only want to hide communication from human users, but also from early detection by computer-aided analysis. Regarding 4), an adaptive audio intrusion detection system is presented in [17].

6 CONCLUSION

Potential attackers can utilize covert physical channels to establish communication between distinct devices that are separated by air gaps and between separated operating system compartments. Attackers might use optical emanations or other parts of the electromagnetic spectrum to establish a covert physical channel. Attackers might also use acoustic (e.g., ultrasonic communication) or even molecular communication for the covert physical channel.

Covert physical channels can be identified by a systematic analysis of potentially cooperating resources within the presented physical environment matrix. The physical environment matrix is targeted at both local operating system and remote covert channels. Designers and evaluators of computing systems might utilize the physical environment matrix to reach a well-grounded decision regarding the absence or presence of a covert physical channel. Finally, different strategies for the mitigation of covert physical channels have been presented. Future work might include more considerations on the application of the physical environment matrix and more details on the mitigation strategies regarding covert physical channels.

REFERENCES

[1] N. Farsad, W. Guo, and A.W. Eckford, “Tabletop Molecular Communication: Text Messages through Chemical Signals,” *PLoS ONE*, Vol. 8, No. 12, Dec. 2013.

[2] B.W. Lampson, “A Note on the Confinement Problem,” *Commun. ACM*, Vol. 16, No. 10, pp. 613–615, Oct. 1973.

[3] M. Hanspach and M. Goetz, “On Covert Acoustical Mesh Networks in Air,” *Journal of Communications*, Vol. 8, No. 11, pp. 758–767, Nov. 2013.

[4] M. Hanspach and J. Keller, “A Taxonomy for Attack Patterns on Information Flows in Component-Based Operating Systems,” in *Proceedings of the 7th Layered Assurance Workshop*, New Orleans, LA, USA, Dec. 2013.

[5] W. van Eck, “Electromagnetic Radiation from Video Display Units: An Eavesdropping Risk?” *Comput. Secur.*, Vol. 4, No. 4, pp. 269–286, Dec. 1985.

[6] J. Loughry and D.A. Umphress, “Information Leakage from Optical Emanations,” *ACM Trans. Inf. Syst. Secur.*, Vol. 5, No. 3, pp. 262–289, Aug. 2002.

[7] D. Genkin, A. Shamir, and E. Tromer, “RSA Key Extraction via Low-Bandwidth Acoustic Cryptanalysis,” http://eprint.iacr.org/2013/857, Accessed: 2013-01-05, Dec. 2013.

[8] M.D. LeMay and J. Tan, “Acoustic Surveillance of Physically Unmodified PCs,” in *Proceedings of the 2006 International Conference on Security & Management*, Jun. 2006.

[9] R. Hasan, N. Saxena, T. Haleviz, S. Zawoad, and D. Rinehart, “Sensing-Enabled Channels for Hard-to-Detect Command and Control of Mobile Devices,” in *Proceedings of the 8th ACM SIGSAC Symposium on Information, Computer and Communications Security,* ser. ASIA CCS’13. New York, NY, USA: ACM, 2013, pp. 469–480.

[10] R. Frankland, “Side Channels, Compromising Emanations and Surveillance: Current and Future Technologies,” Department of Mathematics, Royal Holloway, University of London, Egham, Surrey TW20 0EX, England, Tech. Rep. RHUL-MA-2011-07, Mar. 2011.

[11] R.A. Kemmerer, “Shared Resource Matrix Methodology: An Approach to Identifying Storage and Timing Channels,” *ACM Trans. Comput. Syst.*, Vol. 1, No. 3, pp. 256–277, Aug. 1983.

[12] R. Langner, “To Kill a Centrifuge. A Technical Analysis of What Stuxnets Creators Tried to Achieve,” http://www.langner.com/en/wpcontent/uploads/2013/11/To-kill-a-centrifuge.pdf, Accessed: 2013-12-15, Nov. 2013.

[13] National Computer Security Center, “A Guide to Understanding Covert Channel Analysis of Trusted Systems,” http://www.fas.org/irp/nsa/rainbow/tg030.htm, Accessed: 2013-06-15, Nov. 1993.

[14] S. Gorantla, S. Kadloor, N. Kiyavash, T. Coleman, I. Moskowitz, and M. Kang, “Characterizing the Efficacy of the NRL Network Pump in Mitigating Covert Timing Channels,” *IEEE Transactions on Information Forensics and Security*, Vol. 7, No. 1, pp. 64–75, Feb. 2012.

[15] M. Luby, “LT Codes,” in *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science*. IEEE, Nov. 2002, pp. 271–280.

[16] P.A. Karger, “Multi-Level Security Requirements for Hypervisors,” in *Proceedings of the 21st Annual Computer Security Applications Conference* (December 05–09, 2005). ACSAC. IEEE Computer Society, 2005, pp. 5–9.

[17] M. Hanspach and M. Goetz, “Recent Developments in Covert Acoustical Communications,” in *Proceedings of Sicherheit 2014*, pp. 243–254, Vienna, Austria, Mar. 2014.