

Attacks on Heartbeat-Based Security Using Remote Photoplethysmography

Robert M. Seepers¹, Wenjin Wang², Gerard de Haan², Ioannis Sourdis³ and Christos Strydis¹

¹Dept. of Neuroscience, Erasmus Medical Center, Rotterdam, The Netherlands

²Dept. of Electrical Engineering, Eindhoven University of Technology, Eindhoven, The Netherlands

³Dept. of Computer Science & Engineering, Chalmers University of Technology, Gothenburg, Sweden

Corresponding author: c.strydis@erasmusmc.nl

Abstract—The time interval between consecutive heartbeats (interpulse interval, IPI) has previously been suggested for securing mobile-health (mHealth) solutions. This time interval is known to contain a degree of randomness, permitting the generation of a time- and person-specific identifier. It is commonly assumed that only devices trusted by a person can make physical contact with him/her, and that this physical contact allows each device to generate a similar identifier based on its own cardiac recordings. Under these conditions, the identifiers generated by different trusted devices can facilitate secure authentication. Recently, a wide range of techniques have been proposed for measuring heartbeats remotely, a prominent example of which is remote photoplethysmography (rPPG). These techniques may pose a significant threat to heartbeat-based security, as an adversary may pretend being a trusted device by generating a similar identifier without physical contact, thus bypassing one of the core security conditions. In this paper, we assess the feasibility of such remote attacks using state-of-the-art rPPG methods. Our evaluation shows that rPPG has similar accuracy as contact PPG and, thus, forms a substantial threat to heartbeat-based-security systems that permit trusted devices to obtain their identifiers from contact PPG recordings. Conversely, rPPG cannot obtain an accurate representation of an identifier generated from electrical cardiac signals, making the latter invulnerable to state-of-the-art remote attacks.

I. INTRODUCTION

The time interval between consecutive heartbeats (cardiac interpulse interval, IPI) is a unique biometric feature which may be used to facilitate security in mobile health (mHealth) solutions, such as body-area networks (BANs) or implantable medical devices (IMDs). In contrast to conventional biometrics, which uses a person's unique physiological features to generate a long lasting person-unique identifier [1], the IPI is a time-varying feature that is known to contain a degree of randomness, making it possible to derive a person- and time-unique identifier from it [2]–[4]. This identifier should be difficult to guess, yet, it is possible to obtain (roughly) the same identifier by simultaneously measuring a cardiac signal of the same person [2], [5]. It is commonly assumed that physical contact is required for measuring these cardiac signals and that only devices trusted by a person can make physical contact with him/her. Under these assumptions, two trusted devices may authenticate with each other if their identifiers are *similar enough*, permitting a (small) disparity between them given the noisy nature of biometric data [3], [5]–[7].

While most studies that suggest *heartbeat-based security* (HBBS) for mHealth assume that physical contact is required to generate the same identifier from IPIs, an increasing number of studies suggest that heartbeats may be measured *remotely* using, for example, Doppler radar [8], capacitive coupling [9], optical vibrocardiography [10], ballistocardiography [11], (thermal) imaging [12], [13] and even through detecting minute difference in a person's voice [14]. Such techniques are primarily intended to have a positive impact on healthcare, for example, by monitoring an infant incubator without touching the frail infant [15]. Nevertheless, they may also provide an *adversary* with a tool for compromising heartbeat-based security through generating an identifier which is *similar enough* to those generated by a trusted (on-body) device *without* requiring physical contact. One of the most prominent threats to heartbeat-based security is *remote photoplethysmography* (rPPG) [16], which measures subtle color variations of a human skin surface using a regular RGB camera [17]. These color variations occur due to changes in the blood volume of human tissue (caused by cardiac contractions), which modify the light absorbed (and reflected) by it. Compared to other non-contact heartbeat-monitoring methods, we expect that rPPG is a particularly strong threat to heartbeat-based security as [16]: (i) it relies on relatively cheap equipment, (ii) it is relatively unaffected by environmental noise and (iii) it can be used from a long distance. Moreover, the (RGB) cameras used by rPPG are ubiquitous in today's society (e.g., webcams, laptops or smartphones), which may allow an adversary to launch a *proxy attack*: For example, an adversary may hack a person's laptop, gain access to its camera and, subsequently, launch an rPPG-based remote attack.

In this paper, we examine whether an adversary could compromise an HBBS system through measuring heartbeats remotely using rPPG. To this end, we implement three state-of-the-art rPPG algorithms called *CHROM* [13], *PBV* [18] and *2SR* [19], and evaluate how accurately these algorithms can obtain IPIs compared to a reference contact PPG (cPPG) sensor, which is often considered as a trusted device in related work [2], [4], [5]. We consider a wide range of settings and parameters that are expected to affect the accuracy of rPPG (and thus security) in practice, including subjects with different skin tones, different light sources, types and intensities of movement, video compression and camera-frame rate. Based on our evaluation, we discuss the threat of remote attacks on

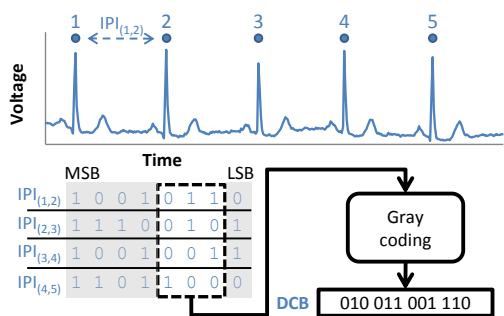


Fig. 1: Generation of a biometric identifier using IPIs, here depicted for an ECG recording.

HBBS and highlight that only identifiers which are generated from electrical cardiac recordings are safe from such attacks. To the best of our knowledge, this is the first study that considers the threat of remote attacks on an HBBS system in detail.

The rest of this paper is structured as follows: We first describe how IPI-based identifiers have commonly been generated and used for authentication in Section II. In Section III, we review related works that detail how similar the identifiers of trusted devices are, as well as studies that consider remote attacks on HBBS. Section IV describes how an adversary could bypass security using rPPG and details the implemented rPPG algorithms. We evaluate how accurately these rPPG algorithms can detect IPIs compared to cPPG in Section V and discuss the practical implications for HBBS in Section VI, after which concluding remarks are provided.

II. BACKGROUND: HEARTBEAT-BASED SECURITY

The IPI is defined as the time difference between two consecutive heartbeats, i.e., $IPI_{(i,i+1)} = beat_{i+1} - beat_i$. It is known that each IPI contains a degree of randomness, which is caused by the balancing action between the sympathetic and parasympathetic nervous systems and which is affected by various (physiological) factors such as smoking, age, gender, emotional and physical state [20], [21], [22]. This randomness makes it possible to derive a unique biometric identifier from IPIs, one prominent way of which is depicted in Figure 1 [2], [3], [5], [23]. First, each device measures a cardiac biosignal (using, for example, electrocardiography (ECG), blood pressure (BP) or PPG), performs peak detection to detect heartbeats (cardiac contractions) and calculates IPIs from consecutive heartbeats. These IPIs are represented as a binary value (in this work assumed to be 8-bit long value), which represents the number of *sampling points* between two heartbeats as measured by a device and, therefore, depends on both: (i) the actual time span between two heartbeats; and (ii) the sampling frequency of a cardiac sensor. From these IPIs, a predefined set of bits is selected: The most-significant bits may be excluded from identifier generation as they are relatively easy to predict, while the least-significant bits may be excluded as they are subject to a high degree of noise, making them difficult to match between different devices. As biometric data is noisy, the selected IPI bits are encoded (using

Gray coding), after which a biometric identifier is formed by concatenating the Gray-coded bits obtained from several IPIs.

If multiple trusted devices simultaneously measure a cardiac signal from the same person, they are expected to generate a similar identifier. A communication protocol may subsequently use these identifiers for authentication purposes [3], [7]. Note that as biometric data is noisy, these identifiers are required to be *similar*, yet, not *identical*. Therefore, authentication is successful if these identifiers are similar enough, i.e., a certain disparity between these identifiers is expected and tolerated.

III. RELATED WORK

An adversary may try to (illegitimately) gain access to an HBBS system by either: (i) generating a biometric identifier (remotely) that is similar enough to an identifier obtained by trusted (on-body) device, as discussed later in Section IV; and/or (ii) finding a flaw in the protocol (or underlying primitives) that uses these identifiers for authentication. While the security of an HBBS system could be compromised through either of these methods, related work has mostly focused on (ii) by both identifying flaws in existing HBBS protocols and proposing new ones [3], [7], [24]. In this paper, we aim to assess the feasibility of (i), i.e., if an adversary may authenticate to an mHealth device illegitimately through remotely obtaining an identifier similar to that of a trusted device. Such an attack could work *regardless* of the underlying protocol and could therefore form a serious threat to HBBS systems in general. Accordingly, we first review the expected disparity between two trusted identifiers (i.e. identifiers obtained by trusted devices). Afterwards, we discuss existing studies that evaluate the feasibility of a remote attack and highlight how our work differs from this related work.

A. Trusted-identifier disparity

Related work describes that the time interval between consecutive heartbeats (IPI) is expected to differ between trusted devices even if they correctly detect these heartbeats. This disparity can be attributed to *inter-sensor variability* (VAR_{is}), that is, the variance between two cardiac recordings due to, for example, the variable pulse-transition time of ventricular contraction to the rest of the body due to pressure differences. The effect VAR_{is} has on the disparity between two devices depends of, among others, where and what type of cardiac signal is recorded and the physiological state of the subject.

Related works have studied the expected disparity due to VAR_{is} for different types of cardiac recordings: 1) *ECG-cPPG*: One identifier is generated from an ECG (chest) and another from a cPPG recording (finger) [2], [5]; 2) *ECG-BP*: One identifier is generated from an ECG (chest) and another from a blood-pressure recording (finger) [2], [5]; and 3) *ECG-ECG*: Both identifiers are obtained from ECG recordings (on the chest) [3], [25]. We present the average bit-error rate (BER) found in each of these studies in Table I (for the 6 least-significant IPI bits). In either model, the BER is highest for the least-significant bits (LSBs) and reduced for more significant bits, i.e., the LSBs are harder to match. Moreover, note that the disparity between two ECG recordings is considerably

TABLE I: Average bit-error rate for different models of $VAR_{i,s}$ between trusted devices.

Dataset	IPI bit #					
	0	1	2	3	4	5
ECG-cPPG [26] (250 Hz*)	–	0.37	0.23	0.15	0.07†	
ECG-BP [27] (250 Hz)	0.46	0.29	0.15	0.08	0.04	0.02
ECG-ECG [3] (360 Hz)	0.08	0.04	0.02	0.01	0.00	0.00

* Derived result from Figure 4 in the paper, resampled from 1000 Hz at 250 Hz.

† BER results for bits 4-5 are reported together in the paper.

lower than an ECG and PPG/BP recording: ECG monitors the electric signal which induces cardiac contraction, whereas PPG and BP depend on changes in blood volume following cardiac contraction. The latter is affected by several physiological phenomena (e.g., pressure differences, the strength of cardiac contractions, etc.), explaining the higher BER.

A second phenomenon that introduces disparity between trusted identifiers is *heartbeat misdetection*. For example, a device may fail to detect a heartbeat (or detect an extra one) due to, among others, movement artifacts. While existing, advanced heartbeat-detection algorithms report a high detection accuracy between 99%-99.9% [28]–[30], it has been demonstrated that even a single undetected heartbeat can cause a significant and unbounded disparity between two identifiers if left unchecked [7], [27]. Nevertheless, heartbeat misdetection can – to some extent – be tolerated through periodic resynchronization during identifier generation [7].

B. Remote attacks

Two existing studies report on the feasibility of remote attacks [3], [31]. In both studies, the probability of a successful remote attack is assessed by placing a 30-FPS (frames per second) webcam at a distance of 50 cm from a subject’s face. This setup closely models the situation in which the subject is working on his/her laptop and where an adversary has gained access to the laptop’s camera (i.e., a *proxy attack*). Despite this similarity, both studies report substantially different results: While Rostami et al. [3] suggests that the 4 LSBs of an IPI can be detected with a 50% accuracy (i.e., no better than guessing the value of an IPI), Calleja et al. [31] describes that over 75% of these LSBs can be detected accurately (compared to an on-body sensor). It should be noted, however, that several crucial details are not provided (e.g., the type of on-body sensor [3] or the used rPPG algorithm [31]), making it difficult to interpret these results. Nevertheless, these studies motivate us to further investigate the feasibility of remote attacks.

In this work, we evaluate the threat of a remote attack in substantially more detail than related works, considering multiple rPPG algorithms and a wide range of parameters and phenomena that are expected to influence these attacks in practice, including skin tone, different light sources, motion and different video formats.

IV. REMOTE ATTACKS USING RPPG

Authentication in HBBS is successful if two identifiers are similar enough. While it is commonly assumed that only trusted devices can generate such similar identifiers, we note

that cardiac contractions induce acute changes in blood volume throughout the body that can be measured remotely (i.e., without physical contact) based on, for example: (i) acute temperature differences (thermal imaging) [12]; (ii) involuntary movement, e.g., head motion (ballistocardiography, Doppler radar) [8], [11]; and (iii) changes in skin color (remote PPG) [13]. Such remote measurements may be used by an adversary to obtain an identifier that is similar enough to an identifier generated by a trusted (on-body) device and, in that way, illegitimately authenticate to trusted devices. In particular, we consider rPPG as a likely candidate for launching a *remote attack* on an HBBS system as, according to a recent survey [16], rPPG – compared to other non-contact heartbeat-monitoring methods: (i) requires relatively cheap equipment (a regular RGB camera) (ii) is relatively unaffected by environmental noise; and (iii) can be used from a long distance. Moreover, the cameras used for deploying such an attack are ubiquitous in today’s society and can be found in, among others, webcams, laptops and smartphones. This may facilitate proxy attacks, in which an adversary obtains access to a target’s own laptop and uses the on-board camera to launch a remote attack.

One of the key challenges for rPPG is to distinguish color differences in an RGB signal due to changes in blood volume (cardiac contractions) from those caused by, among others, luminance or motion. A widely applicable and reliable remote attack would overcome these issues to maximize the probability of breaching security. While several such rPPG algorithms have recently been proposed [32], the evaluation of these algorithms rarely reports on how accurately *individual* heartbeats can be detected, making it difficult to directly assess the feasibility of a remote attacks from existing studies. To this end, we choose to implement and evaluate three state-of-the-art rPPG algorithms¹ which – based on a recently published survey [32] – are expected to have a high detection accuracy:

- *Chrominance (CHROM)* [13]: The CHROM algorithm assumes specular reflection and intensity variations as the common challenges for accurate rPPG. It removes the specular reflection first by projection on the chrominance plane and removes remaining distortions (primarily the intensity variations) in the pulse-signal by a real-time tuning.
- *Blood-volume pulse vector (PBV)* [18]: PBV uses a pre-calculated blood-volume pulse vector (a unit-length color vector) as a prior to compute the weights for extracting the pulse from the RGB color-channels and suppress the (motion-induced) distortions.
- *Spatial Subspace Rotation (2SR)* [19]: The 2SR algorithm obtains a set of skin pixels from its RGB input and transforms it into a spatial subspace. This subspace is subsequently tracked over time to detect changes in hue, where the hue-change is measured as the temporal rotation of the spatial subspace of skin-pixels. Since the pulse-signal is extracted from the hue only, 2SR is inherently independent of all intensity and color-

¹A detailed explanation of these algorithms and the main differences between them can be found in [32].

saturation variations, which particularly benefits the situations where (motion-induced) intensity changes and specular-reflection variations dominate the RGB signal.

V. EVALUATION

A successful remote attack can be mounted if rPPG can obtain IPIs with similar accuracy as a trusted device. We next evaluate this accuracy for a wide range of recording parameters and physiological phenomena, using a cPPG sensor as a model for a trusted device (a choice commonly made in related work [2], [4], [5]). We next detail our experimental setup, after which our evaluation ensues in Subsection V-B.

A. Experimental Setup

In this subsection, we first describe the setup used to create a dataset of rPPG and cPPG signals for a wide range of parameters. Afterwards, we discuss the metrics used to evaluate how accurately rPPG can measure IPIs compared to cPPG. This study has been approved by the Internal Committee Biomedical Experiments of Philips Research and informed consent has been obtained from each subject participating in our experiments.

1) *Dataset*: To evaluate the feasibility of remote attacks, we constructed a dataset containing 71 video sequences and corresponding cPPG recordings. The videos were recorded with a regular RGB video-camera², which provides a spatial resolution of 768×576 pixels, 8-bit depth and a frame rate of 20 FPS, and were stored in an uncompressed bitmap format and constant frame-rate. The cPPG recordings were obtained using a finger-based transmissive pulse oximeter³ and were synchronized with the video frames.

It is expected that a remote attack is most likely successful if a subject is stationary and proximal to the camera, as it minimizes the environmental influences that distort the rPPG signal. We expect that such a setting can occur in practice (e.g., a proxy attack on a patient working from his or her laptop) and, thus, forms a primary concern for the security of an HBBS system. As a default setup we, therefore, place the camera approximately 1 meter in front of the subject at rest (sitting), resulting in approximately 30,000 skin-pixels in each image. The default subject is a male adult with skin-type III according to the Fitzpatrick scale with his face visible on the camera [33]. The subject is illuminated by a frontal fluorescent lamp⁴ and the duration of each recording is around 90 seconds (i.e., 1800 frames by a 20 FPS camera). It is expected that the accuracy at which individual IPIs can be detected using rPPG depends on various physiological phenomena and parameters, such as skin composition (fat), arterial stiffness, skin temperature, skin tone, motion and luminance. While several of these parameters (e.g., arterial stiffness) cannot be controlled or measured in our experiments, we investigate the accuracy at which rPPG can detect heartbeats in a wide range of circumstances by varying the following parameters of our

default setup:

- **Skin tone**: Skin tone directly influences the light reflected from the skin (primarily the diffuse reflections that contain pulsatile information [19]) and is, therefore, expected to affect rPPG signal quality. To investigate if a part of the population is at a greater risk of remote attacks, we record rPPG and cPPG from 15 subjects with various skin tones. These subjects were categorized into three skin-types based on the Fitzpatrick scale: 5 Western European subjects (*skin-type I-II*), 5 Eastern Asian subjects (*skin-type III*), and 5 Sub-Sahara Africa/Southern Asian subjects (*skin-type IV-V*).

- **Light source**: Related work has found that color variations due to blood-volume changes are primarily expressed in the green wavelength [13]. It may become harder (or easier) to identify these color variations if different (types of) light sources are used, hence, affecting the accuracy of rPPG. We, therefore, change the (fluorescent) light source employed in our default setup (using our default test subject with skin-type III) to 6 different sources, including red, green, blue, red-green, red-blue and green-blue LED lamps⁵.

- **Motion**: Motion is a challenging factor for accurate rPPG as it distorts the light reflected from (a particular patch of) skin [19]. In practice, motion occurs in a wide range of circumstances which would allow an adversary to launch an attack on a unknowing subject, for example, a subject walking on the street or sitting on the bus. To evaluate the effect of motion on rPPG accuracy, we record 1 video from our default subject whilst rotating his head continuously (with a rotation angle between 160 and 180 degrees and at a rate of roughly 2 seconds per rotation), as this form of motion was previously found to be a substantial challenge for rPPG [34].

- **Exercise recovery**: To evaluate the robustness of rPPG with respect to pulse-rate changes, a series of videos is recorded to analyze the pulse-rate recovery after exercise. We recruited 6 subjects (3 males and 3 females) of skin-types I-III to participate in this evaluation. Each subject performed 3 different levels of running (with different intensities) by adjusting the speed and gradient of a treadmill: *low* (gradient=12°, speed=4-5 km/h), *medium* (gradient=14°, speed=5-6 km/h), and *high* (gradient=15°, speed=7-8 km/h). The duration of each running exercise is 3 minutes. After the exercise, the subject immediately sits in front of the camera for a recording.

- **Video compression**: It is not uncommon for e.g. webcams to employ (on-board) compression on their raw input signal, resulting in a lower signal quality. As rPPG may be mounted from such commodity hardware, we compress the 15 (uncompressed) videos recorded in the skin-tone category with Motion JPEG encoding (the default implementation of Motion JPEG 2000 in MATLAB), which is commonly employed in commercially available cameras, and evaluate

²Global shutter RGB CCD camera USB UI-2230SE-C from IDS.

³ContecMedical model CMS50E.

⁴Philips HF3319 - EnergyLight White.

⁵Philips LivingColors Bloom.




Fig. 2: Whisker plot depicting the BER (bit-0 to bit-3) of the rPPG algorithms for the entire dataset. Outliers are indicated as red crosses.

the resulting rPPG accuracy [35].

- **Video frame-rate:** While our baseline setup exploits a 20-FPS camera, related work that studies identifiers generated by trusted (contact) devices typically reports on substantially higher sampling rates (≥ 250 Hz) [3], [26], [27]. While such high sampling rates cannot be obtained using the recording equipment available to us, we record 1 video at a frame rate of 50 FPS (the fastest setting possible for our camera) to identify if a higher frame rate can yield an improvement in accuracy.

2) *Evaluation metrics:* A feasible remote attack on our emergency-access scheme requires that individual heartbeats are detected *correctly* and *accurately*. We first measure how many heartbeats are detected *correctly* by rPPG, using cPPG as a point of reference. We report this result in terms of the false-positive rate (FPR) and the false-negative rate (FNR), i.e., the rate at which rPPG incorrectly detects signal artifacts as heartbeats and the rate at which rPPG fails to detect true heartbeats, respectively. Heartbeats in both the rPPG and cPPG recordings were detected using an in-house, validated peak-detection algorithm, after which FNR and FPR were determined. While the cPPG and rPPG recordings were synchronized at the start of their recordings, we allowed some slack between the recordings by considering any two heartbeats, which were detected by both cPPG and rPPG within a range of 25% the camera-frame rate of each other, as a true positive.

After identifying the FNR and FPR, we have obtained a set of heartbeats that are correctly detected by both rPPG and cPPG. We derived IPIs from consecutive heartbeats that were detected correctly and determined how *accurately* rPPG can obtain IPIs compared to cPPG. In line with related work, this accuracy is reported in terms of the average bit-error rate (BER), i.e., the average rate at which a particular rPPG-based IPI bit does not match a cPPG-based one.

B. Experimental Results

We next evaluate how accurately rPPG may obtain IPIs compared to cPPG. We first compare the three rPPG algorithms by considering their detection results for our entire dataset as a whole, after which we discuss the results for individual dataset parameters in more detail.

Figure 2 depicts a statistical comparison of the BER obtained by the different rPPG algorithms for the entire parameter space. The most significant bits (bit 4 onwards) were

TABLE II: FPR, FNR and BER results for various datasets using the 2SR algorithm.

Category	Type	FPR	FNR	BER			
				bit 0	bit 1	bit 2	bit 3
Skin tone	Type I-II	0.02	0.02	0.34	0.02	0.00	0.00
	Type III	0.01	0.01	0.48	0.10	0.00	0.00
	Type IV-V	0.05	0.04	0.47	0.24	0.02	0.00
Light source	Fluorescent	0.00	0.00	0.44	0.01	0.00	0.00
	Various	0.00	0.01	0.50	0.46	0.00	0.00
Body motion	Stationary	0.00	0.00	0.44	0.01	0.00	0.00
	Rotation	0.01	0.01	0.61	0.29	0.00	0.00
Exercise recovery	Low	0.01	0.01	0.47	0.06	0.00	0.00
	Mid	0.02	0.03	0.32	0.72	0.00	0.00
	High	0.05	0.05	0.40	0.59	0.00	0.00
Video compression	uncompr.	0.03	0.02	0.43	0.12	0.01	0.00
	compr.	0.07	0.05	0.46	0.17	0.02	0.00
Frame rate	20 FPS	0.00	0.00	0.44	0.01	0.00	0.00
	50 FPS*	0.02	0.00	0.12	0.00	0.00	0.00

* Rescaled to 25 FPS to be inline with other reported BERs.

always detected correctly (BER = 0.0) and are, therefore, not depicted. For each of the algorithms, we find a high BER ≈ 0.5 for the least-significant IPI bit in an rPPG recording (bit 0) obtained using a 20-FPS camera, i.e., rPPG cannot detect this bit with high accuracy. The BER is gradually reduced when moving to more significant bits, indicating that rPPG can detect these bits with improved accuracy. The globally averaged BER values of bit-0 through bit-3 are: (i) CHROM - 0.43, 0.24, 0.00, 0.00; (ii) PBV - 0.45, 0.30, 0.03, 0.00; and (iii) 2SR - 0.42, 0.21, 0.00, 0.00. Our evaluation suggests that the 2SR algorithm detects IPIs with (slightly) better accuracy than the CHROM and PBV algorithms.

We next consider the effect of various, individual parameters on the detection performance. Table II reports the FNR, FPR and BER for the various categories in our dataset obtained using the 2SR algorithm (the CHROM and PBV algorithms show similar albeit slightly worse results and are, therefore, not discussed in detail). Starting with different *skin tones*, we find that a darker skin makes heartbeats both more difficult to detect accurately (increased BER). This is explained by the higher melanin content in dark skin compared to light skin which reduces the diffuse reflections that contain pulsatile information, lowering signal quality and, accordingly, detection performance. In a similar way, we notice that the detection performance depends on the illumination spectrum. From the various light-sources that we tested, the fluorescent lamp gives the highest performance, while light sources with a less balanced radiation energy in the RGB channels seem to reduce the detection accuracy, even for stationary subjects.

Motion (rotation of the head) affects the accuracy of rPPG as it distorts the amount of light reflected from the skin to the camera. While rotation can be considered as a relatively easy challenge for rPPG (as the motion pattern is regular [19]), it still leads to a substantial increase in BER. Besides, we find that the accuracy of rPPG is greatly reduced when obtained from subjects during *exercise recovery*. These subjects often sweat profusely and breathe heavily which results in unintended body motion and, accordingly, a reduction in detection accuracy (albeit not as substantial as when rotating the head, which yields more severe body motion).

In terms of hardware (video) parameters, we notice that *video compression* yields a substantial increase in terms of

FPR and FNR, yet hardly affects BER. That is, the employed compression sometimes causes signal artifacts that result in incorrectly detected heartbeats, yet it does not noticeably affect the time interval between *correctly detected* heartbeats. Finally, we increase the camera-frame rate to 50 FPS and record one video from our default subject. A fair comparison between this and our other experiments – which employ a 20-FPS camera – requires that the individual IPI bits represent a similar frequency content. To this end, we first halve the IPI values obtained using our 50 FPS experiment prior to calculating the BER, effectively representing the BER values as-if they were obtained using a 25 FPS camera⁶. Note that the BER is considerably lower than when a 20 FPS camera is used, i.e., a higher camera-frame rate can increase the accuracy of rPPG.

VI. DISCUSSION

In the previous Section, we evaluated the accuracy at which rPPG can obtain IPIs for a wide range of factors and demonstrated that rPPG can obtain IPIs with similar accuracy as a cPPG sensor in various cases. A successful remote attack requires that an adversary obtains IPIs with similar accuracy as trusted devices. We may, thus, assess the threat of a remote attack by comparing the accuracy of rPPG (compared to a reference cPPG sensor) to the expected BER for trusted devices (discussed in Section III).

Table III presents the expected BER for an adversary (rPPG-cPPG) and various BER-models for trusted devices. Here, we compare the results obtained using our 50-FPS camera setup as it both: (i) models a remote attack in a realistic setting; and (ii) has the lowest BER among our experiments. That is, these results can be used to determine the (worst case) security performance of a heartbeat-based-security system. Note that the sampling (frame) rates used in our experiments are considerably lower than those reported for trusted (contact) devices. To facilitate a more direct comparison between these results, we again rescale the rPPG results to represent a similar frequency content (in doing so, we effectively shift the BER results of our 50-FPS camera to more-significant bits).

Let us first consider the feasibility of a remote attack on an HBBS system that allows trusted devices to obtain their IPIs using both electrical (ECG) and blood-volume-based cardiac recordings (BP and PPG). Such systems are often designed to tolerate the relatively high BER expected between such recordings (i.e., the *ECG-BP* and *ECG-cPPG* entries in Table III). Related work that proposes to secure mHealth using such systems – and which assumes that an adversary can only *guess* the value of a biometric identifier – often suggests that the most *secure* identifiers are generated from the most-significant IPI bits (from IPI-bit position 3 onwards) [4], [27], [36]. Our comparison in Table III, however, suggests

⁶Recall from Section II that an IPI value is determined by the number of sampling points that can be measured by a sensor between two heartbeats. Halving the sampling frequency means obtaining half the sampling points and, hence, would result in half the IPI value, effectively shifting the BER results by 1. For completeness, the actual BER values obtained using our 50-FPS camera (without rescaling) were 0.44, 0.12, 0.00 and 0.00, respectively, for IPI bits 0 through 3.

that rPPG may detect these IPI bits with similar accuracy as trusted devices, even using the relatively cheap (commodity) hardware employed in our study. In other words, we expect that an adversary can generate an identifier that is similar enough to those employed in such an HBBS systems and, thus, compromise security.

A more constrained HBBS system requires all trusted devices to measure their cardiac signals electrically. This limits which mHealth devices can rely on HBBS for authentication, as electrical cardiac signals (e.g., ECG) may not be detected throughout the human body. Nevertheless, it also lowers the BER significantly as these signals are less substantially affected by noise, as reported for the *ECG-ECG* entry in Table III. This low BER permits such HBBS systems to employ secure identifiers that are formed from the 4 least-significant IPI bits (bits 0-3) [3], [7], [23]. In our study, we were unable to obtain an accurate representation of these IPI bits with rPPG, even under our best measuring conditions (stationary subject with a light skin color, fluorescent-light source and a camera-frame rate of 50 FPS). Moreover, related work has found these bits to be independently distributed, i.e., even if the most-significant IPI bits are captured using rPPG, they cannot be used to estimate the value of an identifier generated from the least-significant IPI bits [23], [36]. Our results, thus, suggest that state-of-the-art rPPG methods cannot provide an adversary with an advantage over merely guessing the value of an identifier, provided that electrical cardiac signals are enforced in the HBBS system. It may be suggested that improving the rPPG algorithm (or otherwise improving the experimental setup by, among others, increasing the camera-frame rate) could further improve the accuracy of rPPG. However, we do not expect that such efforts would significantly increase the threat of remotely attacking these systems, due to the following reasons:

- Commodity hardware often records at a default camera-frame rate of 50 FPS. A more precise rPPG recording would require a higher frame rate, which makes it difficult (or impossible) for an adversary to launch a proxy attack (as it requires overriding the default settings of the targeted subject's camera). As such, the adversary would be required to be *proximal* to his target (subject), which significantly lowers the probability (and, thus threat) of a remote attack;
- Even if an adversary is proximal to his target and can record at an appropriate frame rate, our results suggest that the target has to be relatively stable, favorably illuminated, with a highly pulsatile body-part, like face or palms, exposed. These requirements make it unlikely that an adversary could launch such an rPPG attack in practice; and
- The electric signal representing heartbeats (the “R-peak” in an ECG) contains substantially higher frequency components that are absent in (generic) PPG. Related work has already described that contact PPG cannot measure the least-significant IPI bits with high accuracy (as can be seen for the *ECG-cPPG* entry in Table III). It can be assumed that rPPG is susceptible to stronger noise

TABLE III: Comparison of average BER for various contact-sensor models (ECG-ECG, ECG-BP and ECG-PPG) and rPPG.

Dataset	Freq. (Hz)	IPI bit #							
		0	1	2	3	4	5	6	7
ECG-BP [27]	250	0.46	0.29	0.15	0.08	0.04	0.02	0.01	0.00
ECG-cPPG [26]	250*	–	0.37	0.23	0.15	0.07 [†]	–	–	–
rPPG-cPPG	200*	–	–	0.44	0.12	0.00	0.00	0.00	0.00
ECG-ECG [3]	360	0.08	0.04	0.02	0.01	0.00	0.00	0.00	0.00
rPPG-cPPG	400*	–	–	–	0.44	0.12	0.00	0.00	0.00

* Derived result from Figure 4 in the paper, resampled from 1000 Hz at 250 Hz.

[†] BER results for bits 4-5 are reported together in the paper.

* Rescaled rPPG results to represent a similar frequency content per IPI bit as reported in related work. These results were obtained from a stationary subject with a light skin tone under fluorescent light and recorded using a 50-FPS camera.

influences (e.g., luminance variations) than cPPG. In other words, even in the unlikely case that an adversary *could* launch a remote attack using rPPG, we expect that it may, at best, be comparable to the result obtained with a contact PPG-sensor, which does not suffice to measure the least significant IPI-bits.

A final point of discussion is if our study sufficiently captures all aspects pertinent to a remote attack. While our experimental setup considers a wide range of parameters that affect rPPG quality, we have only considered these in the context of a single subject sitting in front of a camera (effectively modeling a subject sitting in front of his or her laptop). It, thus, remains to be seen if a remote attack could be launched in all possible scenarios: For example, it may be more difficult to correctly detect a subjects IPIs if he or she is moving or is in a crowd. Furthermore, we have so far implicitly assumed that a remote attack can directly be launched on any wearable or implantable device that employs HBBS. Such devices may, however, face additional security measures that can thwart such attacks: Modern implantable cardiac defibrillators (ICDs), for example, require that the ICD is first activated using short-range (< 10 cm) communication before enabling long-range communication. In such cases, an adversary (or an accomplice) would first have to activate the long-range communication by getting in close proximity to a target subject, severely limiting the practicality of a remote attack. While the security of existing implementations of this mechanism is questionable [37], we recommend that future work further explores if such mechanisms could hamper a remote attack.

VII. CONCLUSIONS

In this paper, we evaluated the feasibility of attacking an HBBS system through measuring heartbeats remotely using rPPG. Our evaluation reveals that an adversary may use rPPG to generate a biometric identifier with accuracy similar to trusted devices that obtain their heartbeats from cPPG or BP recordings, allowing an adversary to breach security. Conversely, it is unlikely that rPPG can generate an identifier that is highly similar to those obtained by trusted devices which employ electrical cardiac recordings (e.g., ECG), even under ideal measurement conditions. We, therefore, expect that heartbeat-based security can be considered secure as long as it is strictly based on electrical cardiac recordings.

REFERENCES

- I. Odinaka, P.-H. Lai, A. D. Kaplan, J. A. O’Sullivan, E. J. Sirevaag, and J. W. Rohrbaugh, “Ecg biometric recognition: A comparative analysis,” *IEEE Transactions on Information Forensics and Security*, vol. 7, no. 6, pp. 1812–1824, 2012.
- C. C. Poon, Y.-T. Zhang, and S.-D. Bao, “A novel biometrics method to secure wireless body area sensor networks for telemedicine and m-health,” *IEEE Communications Magazine*, vol. 44, no. 4, pp. 73–81, 2006.
- M. Rostami, A. Juels, and F. Koushanfar, “Heart-to-heart (H2H): authentication for implanted medical devices,” in *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*. ACM, 2013, pp. 1099–1112.
- R. M. Seepers, C. Strydis, I. Sourdis, and C. I. De Zeeuw, “Enhancing Heart-Beat-Based Security for mHealth Applications,” *IEEE Journal of Biomedical and Health Informatics*, vol. pp. no. 1, pp. 1–9, 2015.
- S.-D. Bao, C. C. Poon, Y.-T. Zhang, and L.-F. Shen, “Using the timing information of heartbeats as an entity identifier to secure body sensor network,” *IEEE Transactions on Information Technology in Biomedicine*, vol. 12, no. 6, pp. 772–779, 2008.
- F. Xu, Z. Qin, C. C. Tan, B. Wang, and Q. Li, “IMDGuard: Securing implantable medical devices with the external wearable guardian,” in *INFOCOM, 2011 Proceedings IEEE*. IEEE, 2011, pp. 1862–1870.
- R. M. Seepers, J. H. Weber, Z. Erkin, I. Sourdis, and C. Strydis, “Secure key-exchange protocol for implants using heartbeats,” in *Proceedings of the ACM International Conference on Computing Frontiers*. ACM, 2016, pp. 119–126.
- O. Boric-Lubecke, V. M. Lubecke, I. Mostafanezhad, B.-K. Park, W. Massagram, and B. Jokanovic, “Doppler radar architectures and signal processing for heart rate extraction,” *Mikrotalasna revija*, vol. 15, no. 2, pp. 12–17, 2009.
- A. E. Mahdi and L. Faggion, “Non-contact biopotential sensor for remote human detection,” in *Journal of Physics: Conference Series*, vol. 307, no. 1. IOP Publishing, 2011, p. 012056.
- U. Morbiducci, L. Scalise, M. De Melis, and M. Grigioni, “Optical vibrocardiography: a novel tool for the optical monitoring of cardiac activity,” *Annals of biomedical engineering*, vol. 35, no. 1, pp. 45–58, 2007.
- G. Balakrishnan, F. Durand, and J. Guttag, “Detecting pulse from head motions in video,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2013, pp. 3430–3437.
- S. Y. Chekmenev, A. A. Farag, W. M. Miller, E. A. Essock, and A. Bhatnagar, “Multiresolution approach for noncontact measurements of arterial pulse using thermal imaging,” in *Augmented vision perception in infrared*. Springer, 2009, pp. 87–112.
- G. de Haan and V. Jeanne, “Robust pulse rate from chrominance-based rPPG,” *IEEE Transactions on Biomedical Engineering*, vol. 60, no. 10, pp. 2878–2886, 2013.
- A. Mesleh, D. Skopin, S. Baglikov, and A. Quteishat, “Heart rate extraction from vowel speech signals,” *Journal of computer science and technology*, vol. 27, no. 6, pp. 1243–1251, 2012.
- L. K. Mestha, S. Kyal, B. Xu, L. E. Lewis, and V. Kumar, “Towards continuous monitoring of pulse rate in neonatal intensive care unit with a webcam,” in *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, 2014, pp. 3817–3820.

- [16] J. Kranjec, S. Beguš, G. Geršak, and J. Drnovšek, "Non-contact heart rate and heart rate variability measurements: A review," *Biomedical Signal Processing and Control*, vol. 13, pp. 102–112, 2014.
- [17] W. Verkruijse, L. O. Svaasand, and J. S. Nelson, "Remote plethysmographic imaging using ambient light." *Optics express*, vol. 16, no. 26, pp. 21 434–21 445, 2008.
- [18] G. de Haan and A. Van Leest, "Improved motion robustness of remote-PPG by using the blood volume pulse signature," *Physiological measurement*, vol. 35, no. 9, p. 1913, 2014.
- [19] W. Wang, S. Stuijk, and G. de Haan, "A Novel Algorithm for Remote Photoplethysmography: Spatial Subspace Rotation," *IEEE Transactions on Biomedical Engineering*, 2015.
- [20] U. R. Acharya, K. P. Joseph, N. Kannathal, C. M. Lim, and J. S. Suri, "Heart rate variability: a review," *Medical and Biological Engineering and Computing*, vol. 44, no. 12, pp. 1031–1051, 2006.
- [21] I. Antelmi *et al.*, "Influence of age, gender, body mass index, and functional capacity on heart rate variability in a cohort of subjects without heart disease," *AJC*, vol. 93, no. 3, pp. 381–385, 2004.
- [22] B. M. Appelhans and L. J. Luecken, "Heart rate variability as an index of regulated emotional responding." *Review of general psychology*, vol. 10, no. 3, p. 229, 2006.
- [23] G.-H. Zhang, C. C. Poon, and Y.-T. Zhang, "Analysis of using interpolate intervals to generate 128-bit biometric random binary sequences for securing wireless body sensor networks," *IEEE Transactions on Information Technology in Biomedicine*, vol. 16, no. 1, pp. 176–182, 2012.
- [24] M. Rostami *et al.*, "Balancing Security and Utility in Medical Devices?" *IEEE Design Automation Conference (DAC)*, 2013.
- [25] G. Zheng, G. Fang, R. Shankaran, and M. A. Orgun, "Encryption for Implantable Medical Devices Using Modified One-Time Pads," *Access, IEEE*, vol. 3, pp. 825–836, 2015.
- [26] S.-D. Bao, "A matching performance study on IPI-based entity identifiers for body sensor network security," in *Biomedical Engineering and Informatics (BMEI), 2012 5th International Conference on*. IEEE, 2012, pp. 808–811.
- [27] R. M. Seepers, C. Strydis, P. Peris-Lopez, I. Sourdis, and C. I. De Zeeuw, "Peak misdetection in heart-beat-based security: Characterization and tolerance," in *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, 2014, pp. 5401–5405.
- [28] P. J. M. Fard, M. Moradi, and M. Tajvidi, "A novel approach in R peak detection using Hybrid Complex Wavelet (HCW)," *International Journal of Cardiology*, vol. 124, no. 2, pp. 250–253, 2008.
- [29] A. Ghaffari, H. Golbayani, and M. Ghasemi, "A new mathematical based QRS detector using continuous wavelet transform," *Computers & Electrical Engineering*, vol. 34, no. 2, pp. 81–91, 2008.
- [30] J. P. Madeiro, P. C. Cortez, J. A. Marques, C. R. Seisdedos, and C. R. Sobrinho, "An innovative approach of QRS segmentation based on first-derivative, Hilbert and Wavelet Transforms," *Medical engineering & physics*, vol. 34, no. 9, pp. 1236–1246, 2012.
- [31] A. Calleja, P. Peris-Lopez, and J. E. Tapiador, "Electrical Heart Signals can be Monitored from the Moon: Security Implications for IPI-Based Protocols," in *IFIP International Conference on Information Security Theory and Practice*. Springer, 2015, pp. 36–51.
- [32] W. Wang, B. den Brinker, S. Stuijk, and G. de Haan, "Algorithmic Principles of Remote-PPG," *IEEE Transactions on Biomedical Engineering*, vol. pp, no. 1, pp. 1–12, 2016.
- [33] T. B. Fitzpatrick, "The validity and practicality of sun-reactive skin types I through VI," *Archives of dermatology*, vol. 124, no. 6, pp. 869–871, 1988.
- [34] W. Wang, B. Balmaekers, and G. de Haan, "Quality metric for camera-based pulse rate monitoring in fitness exercise," in *Image Processing (ICIP), 2016 IEEE International Conference on*. IEEE, 2016, pp. 2430–2434.
- [35] S. Hanfland and M. Paul, "Video Format Dependency of PPGI Signals," *International Student Conference on Electrical Engineering (POSTER)*, 2016.
- [36] R. M. Seepers, C. Strydis, I. Sourdis, and C. I. De Zeeuw, "On Using a Von Neumann Extractor in Heart-Beat-Based Security," in *IEEE Trustcom*, 2015, pp. 491–498.
- [37] E. Marin, D. Singelée, F. D. Garcia, T. Chothia, R. Willems, and B. Preneel, "On the (in) security of the latest generation implantable cardiac defibrillators and how to secure them," in *Proceedings of the 32nd Annual Conference on Computer Security Applications*. ACM, 2016, pp. 226–236.