Adaptive Entity-Identifier Generation for IMD Emergency Access

Robert M. Seepers[1], Christos Strydis[1], Ioannis Sourdis[2] and Chris I. De Zeeuw[1]
[1]Dept. of Neuroscience, Erasmus Medical Center, Rotterdam, The Netherlands
[2]Dept. of Computer Science & Engineering, Chalmers University of Technology, Gothenburg, Sweden
{r.seepers, c.strydis, c.dezeeuw}@erasmusmc.nl sourdis@chalmers.se

ABSTRACT

Recent work on wireless Implantable Medical Devices (IMDs) has revealed the need for secure communication in order to prevent data theft and implant abuse by malicious attackers. However, security should not be provided at the cost of patient safety and an IMD should, thus, remain accessible during an emergency regardless of device security. In this paper, we present a novel method of providing IMD emergency access, based on generating Entity Identifiers (EI) using the Inter-Pulse Intervals (IPIs) of heartbeats. We evaluate the current state-of-the-art in EI-generation in terms of security and accessibility for healthy subjects with a wide range of heart rates. Subsequently, we present an adaptive EI-generation algorithm which takes the heart rate into account, maintaining an acceptable emergency-mode activation time (between 5-55.4 s) while improving security by up to 3.4x for high heart rates. Finally, we show that activating emergency mode may consume as little as $0.24\mu J$ from the IMD battery.

Keywords

Inter-pulse interval, security, accessibility, implant, low power

1. INTRODUCTION

It has recently become apparent that several commercially available IMDs do not employ any form of security [5]. Non-secure communication may lead to an adversary stealing private data or, worse, altering device parameters or even disabling the IMD. While an IMD has strict security requirements, it has to remain accessible in case of emergency (when a patient cannot him- or herself authenticate to the IMD) as a patient's safety is more important than device security [6]. For example, an emergency team might require the latest data logs from the IMD or to be able to disable the IMD on short notice and in the field. Moreover, due to the ultra-low-power nature of IMDs, security has to be provided at minimal power and energy consumption.

In this paper we present a novel method of providing emergency-access to IMDs, based on generating Entity Identifiers (EIs) using Inter-Pulse Intervals (IPIs) of heartbeats. IPI-based EI generation may be an attractive solution to providing emergency access for the following reasons: *1)* Heart beats are universally present in humans, allowing IPI-based security to be used by the general population; *2)* a patient cannot "lose" his or her heart beat, providing availability; and *3)* a physical connection is required between the reader and patient for a period of time, reducing the likelyhood of obtaining unauthorized access to the IMD.

The rest of this paper is structured as follows. First, we will briefly discuss works related to IPI-based security in Section 2. We will describe the current state-of-the-art, static, EI-generation algorithm and evaluate it in terms of security and accessibility in Section 3.1. In Section 3.2, we present an adaptive EI-generation algorithm which improves the static algorithm by considering the heart rate (HR), and evaluate its security and accessibility aspects. We estimate the IMD-resource overheads in Section 3.3 and, finally, concluding remarks will be given in Section 4.

2. RELATED WORK

The IPI is the time interval between two consecutive heart beats. It has been shown that each IPI contains a number of bits with a high degree of entropy and, by combining a number of subsequent IPIs, a security key may be generated [9]. As both a reader and an IMD may generate a key from the same cardiac signal, the generated key may be used as an Entity Identifier (EI) for authentication [3].

Previous work has analyzed the security of EIs generated in the context of Body-Area Networks (BANs) using a static EI-generation algorithm. Poon et al. [9] and Bao et al. [3] have evaluated the security of EIs generated from healthy and hypertensive subjects, showing that four entropic bits are present in these subjects, resulting in a correct classification of over 97%. Zhang et al. [13] have analyzed patients with cardio-vascular disorders (CVDs) using a number of National Institute of Standards and Technology (NIST) statistical tests. They have shown that healthy subjects contain five entropic bits per IPI, whereas patients with CVD shown a reduction in entropy of roughly 20%. Moreover, they have shown that IPIs are independently distributed, i.e. the value of an IPI may not be derived from previous IPIs.

In this work, we further analyse the static-EI generation algorithm by including subjects with higher functional capacity (resulting in lower heart rates), which has been shown to increase the entropy in medical literature [2]. Further-

CS2, January 20 2014, Vienna, Austria
Copyright 2014 ACM CS2 978-1-4503-2484-7/14/01 ...$15.00
http://dx.doi.org/10.1145/2556315.2556324.

more, we evaluate subjects during exercise which may show a significant reduction (of over 75%) in entropy [11]. Moreover, we consider the time required to generate an EI within the context of providing emergency access to an IMD and provide an adaptive EI-generation algorithm which significantly improves the security performance of generated EIs.

3. IMPLEMENTATION

As discussed in Section 2, the IPI may be used for generating an EI for authentication. We propose to use this EI to activate an IMD's emergency mode: The EI generated by the reader EI_r may be sent to the IMD and, if it matches the EI generated by the reader (EI_k), the IMD opens its communication channels for emergency access. The EI should be generated in such a way that only authorized readers are able to activate emergency mode and within an acceptable time frame. In the following Sections, we evaluate the accessibility and security performance for both the static and adaptive EI-generation algorithms for subjects with a wide range of heart rates (HRs), after which we estimate the IMD-resource overheads for generating EIs.

3.1 Static EI Generation

In the static-EI generation algorithm [9], both the reader and IMD first collect a number of IPIs from an HR sensor. For each IPI, a fixed number of bits are extracted: Higher bits are discarded as there is little entropy available in these bits, while lower bits are removed as they contain a significant sensor-variation component. As the reader and IMD measure the IPIs from two different locations, slight differences may occur between the IPIs sampled on the reader and IMD, which may result in a disparity between the two EIs even if they are derived from the same signal. Gray coding is applied to increase robustness against these differences, after which the bits are concatenated to form an EI. Finally, the IMD evaluates the Hamming distance between the EIs generated by the reader and IMD and allows access if the Hamming distance is below a certain threshold, i.e. the EIs are similar enough.

3.1.1 *Experimental Setup*

For our experiments, we have employed the following datasets: *1) The Ironman dataset* from the Meditation database [8, 4]. From this dataset, we have used all nine recordings of subjects with a relatively low HR due to increased functional capacity (mean HR: 48.9 Beats-Per-Minute (BPM)); and *2) The Rest-and-Exercise dataset* from the BioSec ECG-database [1]. From this dataset, we have used 58 ECG-recordings from subjects at rest (mean HR: 75.8 BPM) and 52 ECG-recordings from subjects immediately after exercise (mean HR: 97.2 BPM). Similarly to related works, we model the measurements of the reader and the IMD as two different leads in an ECG. As our used datasets contain single ECG-leads only, we have emulated recordings from a second lead using realistic inter-sensor variance obtained from the *Fantasia dataset* [7]. The datasets have been scaled to their lowest common sampling frequency of 200 Hz so as to provide an unbiased comparison. In order to examine the security and accessibility performance of the EI-generation algorithm as a function of HR, we have divided the total dataset in five equidistanced HR bands (30-55, 55-80, 80-105, 105-130 and 130-155 BPM) and evaluate these bands individually.

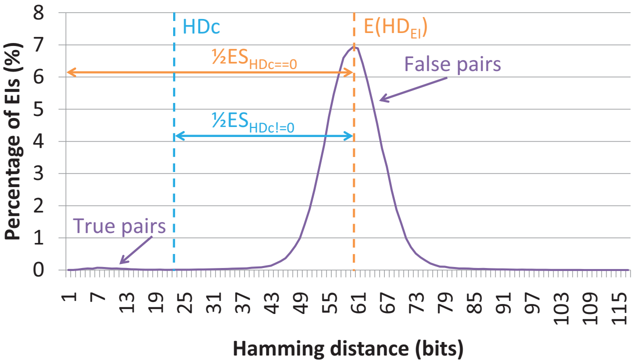

Figure 1: EI-strength approximation for 120-bit EIs and HR band 55-80 BPM using two bits per IPI.

We evaluate the EI-generation algorithm using a fixed EI length and vary the number of bits extracted from each IPI to obtain this EI. We have chosen to use a fixed EI length of 120 bits, as *1)* this is conveniently dividable by numbers one through six, allowing us to test a wide array of settings; and *2)* it may still be considered representable for a 128-bit EI, allowing a more direct comparison to related works which have typically used 128-bit EIs. Note, however, that the EI-generation algorithms discussed in this paper are not limited to a particular EI length.

As safety (accessibility) outweighs security in IMDs, we need to ensure that EIs are generated within an acceptable time frame and a true EI pair (two EIs measured at the same time and from the same person) is never rejected. Without loss of generality, we generate the most secure EI possible under the assumption that emergency mode should be available within 60 *sec*. To evaluate the worst-case access time, we determine the time required to generate a EI using the lowest average HR available in our dataset (37 BPM). In order to make sure that all true EI pairs are accepted, we determine the Hamming distance classifier (HDc) at which no true EI pairs are rejected.

Figure 1 depicts a typical distribution HD_{EI} of Hamming Distances resulting from the all-to-all comparison of EIs within one HR band; the mean value of the distribution is denoted by $E(HD_{EI})$. $E(HD_{EI})$ is used to estimate the available entropy in the generated EIs: Every bit of entropy in an EI results in a 50% chance of matching a bit between two EIs and, accordingly, increases $E(HD_{EI})$ by 0.5. As such, we may approximate the number of entropic bits by $H(EI) = 2 \cdot E(HD_{EI})$. Further depicted in Figure 1 is the Hamming-distance classifier HDc. If HDc equals zero, i.e. both the reader and the IMD contain an exact copy of the same EI, which is common in many cryptosystems, the EI security strength (ES) is equal to the available entropy, i.e., $ES_{HDc=0} = H(EI)$. However, as EI pairs are rarely identical in our case, we are forced to use a non-zero HDc. In effect, this reduces the EI strength and, thus, we approximate the EI strength as $ES = 2 \cdot (E(HD_{EI}) - HDc)$.

3.1.2 *Experimental Results*

Table 1 shows the worst-case access time (WCAT) and EI-classification criterion (Hamming Distance, HDc) required to guarantee that all true EI pairs are accepted, for various numbers of bits used per IPI. Naturally, by increasing the number of bits extracted per IPI, less IPIs are required to generate a EI and the WCAT is reduced accordingly. We, thus, conclude that the accessibility benefits from extracting

Table 1: Accessibility performance for static EI-generation

#bits/IPI	#IPIs	WCAT (s)	HDc
1	120	194.6	39
2	60	97.2	23
3	40	64.9	16
4	30	48.6	14
5	24	38.9	12
6	20	**32.4**	10

more bits per IPI. Moreover, Table 1 shows that **at least four bits should be used per IPI to generate a EI with an access time of <60 s**.

Column 4 of Table 1 reports the HDc at which no true EI pairs are rejected. The HDc starts relatively high when only one bit is extracted per IPI and becomes smaller as more bits are used, implying that true EI pairs become more similar. This is an expected result since *1)* the sensor variation is less noticeable in higher bits, thus causes less disparity in a true EI pair; and *2)* we expect a reduction in entropy in higher bits, leading to more similar EIs (true and false EI pairs alike). Note that the number of bits extracted in Column 1 does not include the lowest three bits of each IPI, as we have found these to contain (too) much inter-sensor variation.

We investigate the available entropy of each HR band by performing an all-to-all EI comparison within each band, as depicted in Figure 2. We can see that, in accordance to related works, the entropy is reduced when more bits are extracted per IPI. Moreover, for a fixed number of four bits per IPI (in order to generate EIs within an acceptable amount of time), we see that **the available entropy in high HRs is decreased by up to 65% compared to low HRs**.

Figure 3 depicts the resulting EI strength for each HR band and number of bits extracted per IPI. We can see that the EI strength is not decreasing monotonically: as the HDc is reduced when more bits are extracted (Column 4 of Table 1), the reduced entropy available may be compensated for. The EI strength in the slowest HR band (30-55 BPM) may be as high as 86 bits when four bits are used per IPI, whereas the highest HRs (130-155 BPM) may provide at most 34 bits of security when only 1 bit is used per IPI. Moreover, in agreement with the entropy available, **a slower HR will generate a more secure EI when a fixed number of bits is used per IPI, up to 8.2x when using four bits per IPI**.

In summary, we have shown that the worst-case access time is too high when *less* than four bits are used per IPI, while using *four or more* bits per IPI may result in a low EI strength. We, thus, conclude that **using a fixed number of bits results in either insufficient accessibility for low HRs or a low degree of security for high HRs**.

3.2 Adaptive EI-Generation

In order to improve the security while maintaining an acceptable access time compared to using a fixed number of IPIs, we propose an algorithm which adapts the number of bits extracted per each IPI, based on the average HR. In so doing, the number of bits extracted per IPI may be increased during slow HRs, maintaining adequate accessibility and EI strength, whereas it may be decreased during high HRs, increasing security while providing suitable access times. This adaptive EI-generation algoritm collects IPIs until it has collected enough EIs for a certain HR band. The average HR is calculated and an EI is generated when enough IPIs have been collected for the corresponding HR band. If not, the algorithm continues to collect IPIs until an EI may be generated (for another HR band).

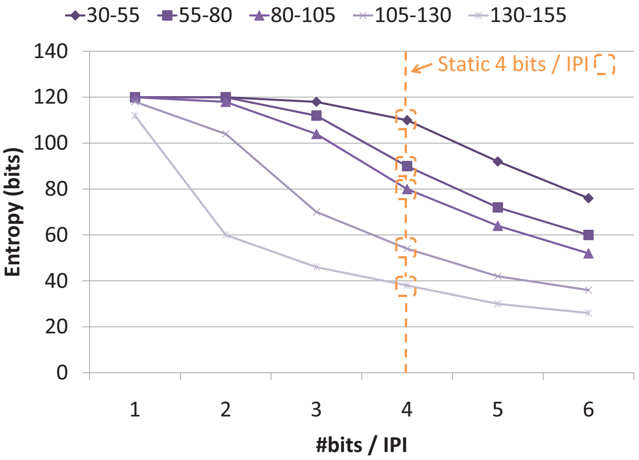

Figure 2: Shannon entropy E(EI) for all HR bands. Also denoted is the selection of #bits/IPI for the static EI-generation case under accessibility constraints.

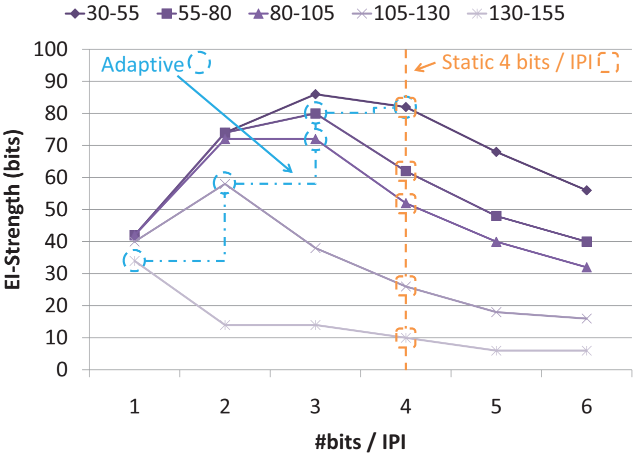

Figure 3: EI-Strength (ES) for all HR bands. Also denoted is the selection of #bits/IPI both for the static and the adaptive EI-generation cases under accessibility constraints.

3.2.1 Security Evaluation

The optimal number of bits used per IPI is depicted in Figure 3 and presented in the second Column of Table 2 for each HR band. As our adaptive EI-generation algorithm uses a number of HR bands, each with a minimal HR, we compute the WCAT for each band as presented in the fourth Column of Table 2. Note that Table 2 includes an HR band of zero BPM: In case of cardiac arrest – which may be detected within 5 *s* – the IMD will recognize the patient is experiencing a cardiac emergency and automatically provide emergency access to the reader. Furthermore, in compliance with our requirements, the access time remains under 60 *s* for all HRs (max: 55.4 *s*). Moreover, we can see that a higher HR may generate an EI within an acceptable timespan using less bits per IPI than a slower HR. Thus, a higher HR (IPIs/timespan) compensates for the increased number of IPIs required.

Table 2: Security and accessibility performance of the adaptive EI-generation algorithm

HR band (BPM)	bits/IPI	#IPIs	WCAT	ES (bits)	ES (%)
0	0	0	5	-	-
30 - 55	4	30	48.6	82	+0
55 - 80	3	40	43.6	80	+29
80 - 105	3	40	30	72	+38
105 - 130	2	60	35.3	58	+123
130 - 155	1	120	55.4	34	+340

Table 3: Execution time on IMD

Function	Ex. time (μs)			
Bits per IPI:	1	2	3	4
EI Generation	385.6	191.9	127.2	95.0

We present the EI strength (ES (bits)) and relative EI-strength increase compared to using four bits per IPI (ES (%)) in the last two Columns of Table 2, respectively. We see that the EI strengths are increased compared to static EI-generation by up to 3.4x for the highest HR band, which reveals to the value of the adaptive EI-gen method. However, it is interesting to observe that the EI strength for the highest HRs is still very low as the 120-bit long EI provides only 34 bits of security. In summary, **our adaptive EI-generation algorithm yields more secure EIs within an acceptable time frame than the static EI-generation algorithm**.

3.3 Resource-utilization overheads

As IMDs are ultra-low-power devices, we have to minimize the resource utilization overheads on the IMD. We estimate the energy and power consumption overheads by profiling the EI-generation algorithm on the Smart-Implantable Security Core (SISC), a 5-stage RISC processor with a 32-bit datapath and a 16-bit ISA with extensions for security applications [10]. We determine the execution time of the EI-generation algorithm in Modelsim and determine the energy consumption by considering the typical, average SISC power consumption. Here, we assume that the IMD readily has access to IPIs, which is a reasonable assumption for e.g. pacemakers. Moreover, modern implantable HR sensors are being developed which can be powered from bio-harvested energy [12] and, thus, cause no additional overhead to the IMD power or energy budgets. Furthermore, we are assuming a protocol is used in which the IMD only replies to the reader when the reader is properly authenticated, i.e., no energy is used on communicating to the reader pre-authentication, whichs limits the energy overhead.

Table 3 shows the execution time of the EI-generation algorithm when 1-4 bits are extracted per IPI. It is shown that when more bits are extracted per IPI, i.e. less samples are required for EI generation, the execution time of the algorithm is reduced. The memory required to store all variables amounts to 500 bytes, which easily fits into the SISC memory. As the SISC core and memories have been measured to consume less than $600\mu W$ of power, we can see that the IMD will consume **no more than $0.24\mu J$ per transaction**, which may be harvested through an IMD wireless RF link, as demonstrated in [10].

4. CONCLUSIONS

In this paper, we have presented a novel method of providing IMD-emergency access, based on authentication through IPI-based EIs. We have shown that the current state-of-the-art in IPI-based EI generation is not capable of providing both a high degree of security for subjects with high HRs and accessibility for subjects with low HRs. To improve the security while guaranteeing high accessibility, we have proposed an adaptive EI-generation algorithm which takes the HR into account, maintaining an acceptable emergency-mode activation time (between 5-55.4 s) while improving security by up to 3.4x for high HRs. Finally, we have shown that activating emergency mode may consume no more than $0.24\mu J$ per transaction.

5. ACKNOWLEDGEMENTS

This work has been supported by the EU-funded project DeSyRe (Grant agreement no: 287611) and would not have been complete without the invaluable feedback of dr. Pedro Peris-Lopez and dr. Foteini Agrafioti.

6. REFERENCES

[1] F. Agrafioti, F. M. Bui, and D. Hatzinakos. Medical biometrics in mobile health monitoring. *Security and Communication Networks*, 4(5):525–539, 2011.

[2] 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*, 93(3):381–385, 2004.

[3] S.-D. Bao et al. Using the timing information of heartbeats as an entity identifier to secure body sensor network. *T-ITB, pp. 772-779*, 12(6), 2008.

[4] A. L. Goldberger et al. Physiobank, physiotoolkit, and physionet components of a new research resource for complex physiologic signals. *Circulation*, 101(23):e215–e220, 2000.

[5] D. Halperin et al. Pacemakers and Implantable Cardiac Defibrillators: Software Radio Attacks and Zero-Power Defenses. In *SP*. IEEE, 2008.

[6] D. Halperin et al. Security and Privacy for Implantable Medical Devices. *PC*, pages 30–39, 2008.

[7] N. Iyengar et al. Age-related alterations in the fractal scaling of cardiac interbeat interval dynamics. *AJP-Regu*, 271(4):R1078–R1084, 1996.

[8] C.-K. Peng et al. Exaggerated heart rate oscillations during two meditation techniques. *IJC*, pages 101–107, 1999.

[9] C. C. Poon et al. A novel biometrics method to secure wireless body area sensor networks for telemedicine and m-health. *IEEE Commun. Mag.*, pages 73–81, 2006.

[10] C. Strydis et al. A system architecture, processor and communication protocol for secure implants. *TACO*, 2013.

[11] M. P. Tulppo et al. Quantitative beat-to-beat analysis of heart rate dynamics during exercise. *AJP Heart*, 271(1):H244–H252, 1996.

[12] R.-F. Xue et al. Ultra-low-power wireless implantable blood flow sensing microsystem for vascular graft applications. In *ISIC*, pages 224–229. IEEE, 2011.

[13] G.-H. Zhang et al. Analysis of using interpulse intervals to generate 128-bit biometric random binary sequences for securing wireless body sensor networks. *T-ITB*, 16(1):176–182, 2012.