Biometric-based Two-level Secure Access Control for Implantable Medical Devices during Emergencies

Xiali Hei and Xiaojiang Du
Department of Computer and Information Sciences
Temple University
Philadelphia, PA 19122, USA
Email: {xiali.hei, dux}@temple.edu

***Abstract*—Implantable Medical Devices (IMDs) are widely used to treat chronic diseases. Nowadays, many IMDs can wirelessly communicate with an outside programmer (reader). However, the wireless access also introduces security concerns. An attacker may get an IMD reader and gain access to a patient's IMD. IMD security is an important issue since attacks on IMDs may directly harm the patient. A number of research groups have studied IMD security issues when the patient is in non-emergency situations. However, these security schemes usually require the patient's participation, and they may not work during emergencies (e.g., when the patient is in comma) for various reasons. In this paper, we propose a light-weight secure access control scheme for IMDs during emergencies. Our scheme utilizes patient's biometric information to prevent unauthorized access to IMDs. The scheme consists of two levels: level 1 employs some basic biometric information of the patient and it is light-weight; level 2 utilizes patients' iris data for authentication and it is very effective. In this research, we also make contributions in human iris verification: we discover that it is possible to perform iris verification by comparing partial iris data rather than the entire iris data. This significantly reduces the overhead of iris verification, which is critical for resource-limited IMDs. We evaluate the performance of our schemes by using real iris data sets. Our experimental results show that the secure access control scheme is very effective and has small overhead (hence feasible for IMDs). Specifically, the false acceptance rate (FAR) and false rejection rate (FRR) of our secure access control scheme are close to 0.000% with suitable threshold, and the memory and computation overheads are acceptable. Our analysis shows that the secure access control scheme reduces computation overhead by an average of 58%.**

***Index Terms*—implantable medical devices; biometric-based security; access control; iris**

I. Introduction

Implantable Medical Devices (IMDs) have been widely used to treat chronic diseases such as cardiac arrhythmia and diabetes. Many IMDs are enabled with wireless communication capabilities and can wirelessly communicate with an outside programmer/reader. With the rapid growth of IMDs, IMD security becomes a critical issue since attacks on IMDs may directly harm the patient. There are a number of attacks that an adversary could launch on IMDs. For example, pacemakers and Implantable Cardioverter Defibrillators (ICDs) contain a magnetic switch (or sensor) that can be activated by sufficiently powerful magnetic fields [1]. Vulnerabilities in the communication interface of wireless programmable IMDs may allow attackers to monitor and alter the function of medical devices without even being in close proximity to the patient [2]. IMDs contain sensitive patient data and information. An attacker could easily launch eavesdropping attacks on IMDs and harvest patient's privacy information using a mobile phone with IMD reader function. Insurance companies also have motivations to perform such passive attacks.

Traditional security schemes (those are designed for sensor networks and other systems) cannot be directly applied to IMDs, due to the severe resource constraints of IMDs, in terms of energy supply, processing, and storage. For example, an IMD manufactured in 2002 (still being used today) contains as less as 8 KB storage [3]. Furthermore, it is not easy to replace the battery for most IMDs, since an IMD is embedded in a human's body and may need a surgery to do so. Hence, it is challenging yet critical to design effective and resource-efficient security and privacy schemes for IMDs.

An intuitive approach for IMD access control during emergencies is to pre-configure a backdoor key in IMDs. In case of emergency, first the medical personnel need to obtain the backdoor key, and then use the key to access the IMD. However, the backdoor-key-based approaches have limitations. Some papers propose to store a global backdoor key in a server, and medical personnel could obtain the key via the Internet. This does not work if the unconscious patient is in another country where the doctors there do not have access to the server. Neither does storing the key in a hospital server. Maintaining a globally available backdoor key is costly.

To sum up, none of the existing IMD access control schemes work well during emergencies. In this paper, we present a novel Biometric-Based two-level Secure Access Control (BBS-AC) scheme for IMDs when the patient is in emergency situations (e.g., in a comma). Most IMDs are embedded in (or closely attached to) a human's body. Based on this fact, we propose novel access control schemes for IMDs by utilizing the human factor. Our BBS-AC scheme has two levels. The first level uses some patient basic biometric information,

including fingerprints' pattern, height, and eye color. The first level provides fast authentication that can defend attackers who do not possess much biometric information of the patient. If an attacker passes the first level, he/she still needs to pass the second level authentication, which uses patient iris image. In our scheme, a clinical personnel does not need to know a key or get a token in advance, and it is not necessary to keep back door keys on a global database. What the clinical personnel need is just devices (e.g., a camera) to obtain a patient's iris image and some basic biometrics.

During emergencies, a patient (say Bob) may be out of conscious and he cannot give his credentials (such as a token and a key) to the medical personnel, nor can he show his ID to tell the doctor about his medical information. Our main idea is to pre-configure a key based on patient's biometrics in an IMD. With a secure access control scheme, a doctor (may be in another country) will be able to access Bob's IMD, obtain his identity and medical information from the IMD, and perform corresponding medical treatments. With the protection from the secure access control scheme, an attacker will not be able to obtain any useful information from the IMD and do harm to the patient.

Biometrics is the technology of recognition or verification of a person's identity by using unique human physical characteristics such as fingerprints, hand geometry, iris, and voice. It provides effective ways for identifications, which can be used for access control and various security functions. Biometrics is better than password/PIN or smart cards thanks to the following traits: no need to memorize passwords; requires physical presence of the person to be identified; cannot be borrowed, stolen, or forgotten. Hence, it is suitable for the IMD access during emergencies. Iris recognition is one of the most precise biometric authentication methods. It is also very fast. Iris is considered as an internal organ, which is protected by eyelid and cornea. In general, it is not easy for an attacker to get a high-quality iris image of a patient. The best way to get a detail-rich iris image is to use a near infra-red (NIR) camera. However, a NIR camera only works well when it is in a distance of 50-70 cm to the person from the front. Within this range, a patient can easily detect a malicious attacker. Due to the above reasons, we choose patient's iris for effective authentications in our access control scheme.

We summarize our major contributions as follows:

1. Based on real iris data sets, we discovered a special bit set in iris code - the Discriminative Bit Set.

2. Via experiments on real iris data sets, we demonstrated that iris recognition can be accomplished by comparing only the Discriminative Bit Set (instead of the entire iris code), which reduces the computation overhead by an average of 58%.

3. We designed a novel secure access control scheme for IMDs during emergencies. The scheme is very effective and has small overhead (suitable for IMDs).

The rest of the paper is organized as follows. In Section II we review the related work on IMD security. In Section III, we present our Biometric-Based two-level Secure Access Control (BBS-AC) scheme for IMDs during emergencies. In Section IV, we discuss the details and results of our performance evaluation. In Section IV, we analyze performance of the BBS-AC scheme. We conclude our work in Section VI.

II. Related Work

A number of literatures (e.g., [3]–[10]) have studied IMDs security issues during non-emergency situations. Some propose to use an (additional) external device such as an access token [4] or a communication cloaker [5]. However, the external devices may be stolen, lost, or forgotten. Certificate-based approaches [6] require the IMD reader be able to access the Internet and a global certification authority needs to be maintained. In [11], the authors propose to let IMD emit an alert signal (sound, vibrations, etc.) when it is engaging an interaction. However, this approach may not work well in noisy environments and it consumes additional battery power of IMDs. Some papers (e.g., [7]–[9]) propose schemes that deny long distance wireless interactions with an IMD unless the proximity of the reader is verified. For example, the secure telemetric link solution in [7] proposes to use a physical backdoor to verify if the reader is close to the IMD. Access control based on close-range communication is very intuitive, however, it is not secure against an attacker that uses special equipment (e.g., high-gain antennas), and it can not prevent the resource depletion attacks [10]. The authors in [9] propose a new IMD access control scheme based on ultrasonic distance bounding. Halperin et al. [3] propose using RF power harvesting for authentications and key exchange to protect IMDs. Hei et al. [10] propose to detect resource depletion attacks by utilizing the patient's IMD access pattern. Although the scheme in [9] could be used for IMD access control when the patient is in emergency, it is difficult to integrate the circuit of the audio receiver into the circuit of IMDs. To sum up, all the above works considered IMD access control when the patient is in non-emergency situations. In this paper, we design a novel access control scheme for IMDs when the patient is in emergencies. Regarding to iris recognition technique, the authors of [17] first presented experiments documenting that some bits in an iris code are more consistent than others.

III. The Biometric-Based Two-Level Secure Access Control Scheme

Many IMDs manufactured today have incorporated certain security functions. Hence, it is reasonable to assume that an IMD has basic security protections. For example, during non-emergency situations, an IMD reader still needs to pass an authentication process in order to access an IMD (e.g., reading data from the IMD). That is, all information (including patient's iris data) that is pre-loaded in an IMD is protected.

In this research, we design a novel Biometric-Based two-level Secure Access Control (BBS-AC) scheme for IMDs during emergencies. The first level employs some basic biometric information of a patient and it is discussed in subsection A. The second level utilizes patient iris images for access control and it is discussed in subsection B.

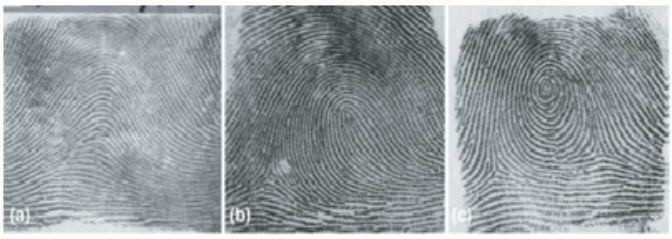

Fig. 1. Three fingerprint types: (a) arch (b) loop (c) whorl [16]

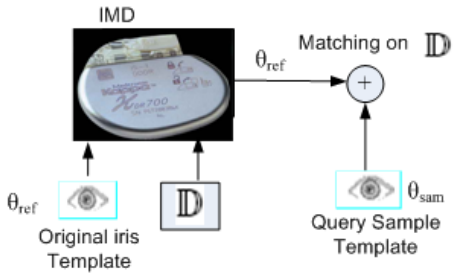

Fig. 2. The iris verification process

A. Level-One Access Control - Using Basic Biometrics

Human's fingerprint pattern is governed by the shape, size, and placement of volar pads [12]. Fig. 1 shows the three fingerprint types: arch, loop, and whorl. Higher and symmetric volar pads tend to form whorls, flatter and symmetric volar pads tend to form arches, while asymmetric volar pads tend to form loops. The fingerprint types of a human does not change in his or her life. Hence, we choose the fingerprint types as one of the basic biometrics in level-one. Level one uses the following three kinds of basic biometrics for access control:

- Fingerprint types of a patient's ten fingers: 0 for arch; 1 for loop; and 2 for whorl.
- Patient's iris color: 0 for dark brown; 1 for light brown; 2 for blue; and 3 for green [13].
- Patient's height.

The three kinds of basic biometrics are stable for an adult. It is not easy for an outside attacker (who does not know the patient) to get all the three kinds of biometric information. Hence, level-one scheme should be able to defend many "random" attackers. The benefits of level-one scheme include:

- The storage required for the basic biometrics is small.
- The computation to verify the basic biometrics is light, which also means the power consumption is low.
- The verification can be done quickly (due to the above two benefits).

Furthermore, level-one scheme gives authorized person (e.g., clinical personnel) easy and fast access to an IMD during a patient's emergency, because the clinical personnel have direct/close contact with the patient and they can easily get the three kinds of basic biometrics. Some extreme cases could prevent medical personnel from obtaining such biometric information. For example, the patient's fingerprints have been destroyed by a fire. In such kind of extreme cases, it is probably more important to perform other medical treatments, rather than try to access the patient's IMD. In this paper, we do not consider such extreme cases. This issue is our future work.

If an attacker is able to collect all the three kinds of biometrics of a patient, then he still needs to pass the level-two authentication in order to access the IMD. We discuss the details of level-two access control in the next subsection.

B. Level-Two Access Control - A New Iris Verification Scheme

1) Obtaining Iris Images: Areas of an iris that are obscured by eyelids, eyelashes, or reflections from eyeglasses, or that have low contrast or a low signal-to-noise ratio lead to errors. In case of emergency, the clinical staff can remove the contact lenses and carry out strict constrains such as no eyelid and no eyelash to shadow iris during the iris image acquisition. Thus, the error from eyelid and eyelash can be significantly reduced.

When configuring an IMD, the highest quality iris image of a patient is chosen as the reference image, and converted into an iris code (a fixed-length binary code). The iris code is pre-loaded in an IMD. During a verification, a sample iris image is obtained and converted into an iris code, which is then used for comparison with the reference iris code. IMD has very limited resource. Hence, in our design, we try to reduce the required computation for iris verification.

2) Generating Iris Codes: The patient's reference iris code θ_{ref} is stored in the IMD before the IMD is implanted into the patient's body. When clinical personnel input a sample iris code θ_{sam}, the IMD compares θ_{sam} with θ_{ref}. In our work, we use the schemes in [14] to generate iris codes from iris images. The iris code length is 9600 bits. This iris verification process is illustrated in Fig. 2.

3) The Discriminative Bit Set of Iris Codes: As mentioned in subsection III.B.1, the noise of iris codes mainly comes from areas that are obscured by eyelids and eyelashes. We focus on these two causes to reduce the noise of iris codes, which increases the accuracy of iris verification. In our research, we wondered whether multiple iris codes of the same eye have some common patterns. We then tried to find patterns among iris codes via experiments on several real iris data sets, including CASIA V1.0 and CASIA-IrisV3-Interval [15]. Fortunately, we were able to find a special bit set among multiple iris codes of the same eye. We refer to this kind special bit set as the Discriminative Bit Set. Based on our experiments on real iris data, we showed that it is possible to perform iris verification by using only a small portion of the iris codes. This greatly reduces the storage and computation requirements of iris verifications, which is significant for resource-limited IMDs.

Discriminative Bit Set - In an iris data set, multiple images are obtained for each iris. For every iris, we choose the clearest image (denotes as image 1) as the reference image. An iris code is generated from each iris image. Recall that an iris code has a fixed length of binary bits (0 or 1). Then we compare the iris code generated from image 1 (the reference code) with iris

codes generated from other images of the same iris. We record the locations of the same bits (denotes as locations12 D_{12}) between the reference code and another iris code 2. Similarly, we record the same locations13 D_{13} between the reference code and iris code 3. We do this for all codes those generated from the same iris. Suppose there are a total of k codes for the same iris. At the end, we obtain the intersection of D_{12}, D_{13},..., and D_{1k}. The intersection includes the common bits of all the iris codes, and it is denoted as set $\mathbb{D}$. The formal definition is given below.

Definition : $D_{ij} = \{d_{ij} |\ d_{ij}\ is\ a\ location\ where\ iris\ code\ i\ and\ j\ have\ the\ same\ bits\}$

$$\mathbb{D} = D_{12} \cap D_{13} \cap ... \cap D_{1k} \quad \text{(III-1)}$$

After we obtain the Discriminative Bit Set $\mathbb{D}$ of each eye's iris, we can use $\mathbb{D}$ for iris code comparison.

We performed experimental study on two real iris data sets (CASIA V1.0 and CASIA-IrisV3-Interval), and recorded the Discriminative Bit Set $\mathbb{D}$ of multiple iris codes for each iris. Furthermore, we did iris-verification tests by using both the Discriminative Bit Set $\mathbb{D}$ and the entire-length iris codes. Our study shows that using only $\mathbb{D}$ for iris verifications provides similar accuracy as that of using the entire-length iris codes. However, using only $\mathbb{D}$ reduces about 58% of the computation overhead of iris verifications. We analyzed a real iris image data set CASIA V1.0 and CASIA-IrisV3-Interval, and obtained the length information of the Discriminative Bit Set $\mathbb{D}$. On average, the ratio of $\mathbb{D}$ to the complete iris code length is about 42%. This shows that using $\mathbb{D}$ for iris verification can significantly reduce storage and computation overheads.

C. The Matching Scheme for Iris Codes

Most iris matching schemes (e.g., the one in [13]) need to compare the entire iris code. In this paper, we propose a novel iris matching scheme, which only uses part of an iris code (Discriminative Bit Sets $\mathbb{D}$).

The Hamming distance is commonly used as a matching metric. In the iris verification case, a Hamming distance gives a measure of how many bits are different between two iris codes. If the Hamming distance is less than the threshold, the verification is considered as successful. Otherwise, the verification fails. There is a Hamming distance threshold (denoted as Th) for iris verification. If the Hamming distance (Hd) is greater than (or equal to) Th, then the two iris codes are considered generated from different irises, and the matching fails. If Hd is less than Th, the two iris codes are considered generated from the same iris.

In our work, we also performed experiments on iris codes by considering noises. Denote the Hamming distances (when considering noises) as Hd. The formal definition is given below:

$$Hd = \frac{1}{m - \sum_{j \in \mathbb{D}} \theta_{refn_j}(OR)\theta_{samn_j}} \sum_{j \in \mathbb{D}} \theta_{ref_j} \bigoplus \theta_{sam_j} (AND)\theta_{refn'_j}(AND)\theta_{samn'_j}. \quad \text{(III-2)}$$

Notes: m is the cardinality of set $\mathbb{D}$, θ_{ref} is the patient's reference iris code, which is pre-stored in the IMD. θ_{sam} is the input sample iris code, which obtained before accessing the IMD and will be verified against θ_{ref}. In the above equation, j belongs to subset $\mathbb{D}$. This means that only bits of subset $\mathbb{D}$ are compared. θ_{refn} and θ_{samn} are the corresponding noise masks (the set of noise bits) for iris code θ_{ref} and θ_{sam}, respectively. And $\theta_{refn'}$ and $\theta_{samn'}$ are the complementary set of θ_{refn} and θ_{samn}, respectively. Compared with the iris matching scheme in [14], our scheme greatly reduces the computation overhead.

IV. Performance Evaluation

A. Experimental Data Sets

In our research, we used real iris image data sets to evaluate our scheme. The iris data sets were collected by the Chinese Academy of Sciences' Institute of Automation (CASIA). We used two data sets - CASIA V1.0 and CASIA-IrisV3-Interval. In our experiments, we used part of the iris images from the two data sets. Specifically, we used a subset V1 (with 264 images) of the CASIA V1.0 data set, and a subset V3 (with 1,370 images) of the CASIA-IrisV3-Interval data set. The total number of iris images that we used is $1,634$, and they are generated from 198 subjects (human). For each iris, we choose the clearest iris image as the reference image, and other images of the same iris are used as training or testing data.

With these iris images, we used the algorithm in [14] to generate iris codes and the corresponding noise masks. Then we performed various experiments by using the iris codes (and noise masks for some tests). The parameters that we chose is the same as those in [14], which generates an iris code of 9,600 bits.

B. Experimental Results

In this subsection, we present our experimental results on iris verification/matching by using our scheme. In our experiments, we use false acceptance rate (FAR) and false rejection rate (FRR) as performance metrics. There is a trade-off between the two metrics. For an iris verification system, if a large Hamming-distance threshold (e.g., Th) is used, then few impostor can fool the system, but many legitimate users would be rejected too. On the other hand, using a small threshold, it rejects less legitimate users, but it also gives more opportunities for a hacker to break into the system. For IMD access control, FRR may outweigh FAR because patient safety outweighs security during emergencies. It would be very costly if a legitimate user (e.g., a doctor) is denied of access to an IMD when a patient has an emergency.

Table I and II list the experimental results of FAR and FRR for different thresholds based on V3 and V1 data sets with noise. As we can see from the tables, if a suitable threshold Th is selected, both FAR and FRR can be 0.000%. The results show that our scheme is very effective in iris verification.

TABLE I
FALSE ACCEPT AND FALSE REJECT RATES WITH DIFFERENT THRESHOLD BASED ON V3 DATA SETS(7 TRAIN IMAGES)

Th	FAR(%)	FRR(%)
0.4	1.254	0.000
0.3	0.000	0.000
0.2	0.000	6.634

TABLE II
FALSE ACCEPT AND FALSE REJECT RATES WITH DIFFERENT THRESHOLD BASED ON V1 DATA SETS (3 TRAIN IMAGES)

Th	FAR(%)	FRR(%)
0.25	18.639	0.000
0.20	0.000	0.000
0.10	0.000	8.728

V. PERFORMANCE ANALYSIS

A. Computation Overhead

It takes *less than 1 ms* to match a pair of iris codes by running Matlab on a computer with a 2.26 GHZ CPU and 3 GB memory. The computation overhead of our matching scheme is only 42% of the scheme in [14]. This kind of saving is significant for resource-limited devices, such as IMDs. Further more, the verification time of our scheme is short, which is critical during emergencies. The level-one authentication only needs to compare (10+2+8)=20 bits, while the level-two authentication need to compare about 9,600*0.42= 4,032 bits. The above discussions show that our level-one scheme is very light-weight, and our level-two scheme significantly reduces the computation overhead (and hence the energy consumption) of iris verification.

B. Storage Requirement

The iris code that we used has a length of 9,600 bits. In addition, the bit set $\mathbb{D}$ needs to be stored in an IMD, and the length of $\mathbb{D}$ roughly equals to 42% of the length of a iris code. Adding the storage of the basic biometric information, the total storage space needed in an IMD is 1,680 bytes. This is a reasonable storage requirement for many current IMDs.

C. Energy Consumption

A pacemaker (with a CPU of type 230) runs at 50 MHZ [18]. And (4,032+20) comparisons only takes about 0.08 ms. The energy consumed is negligible compared to ordinary therapy or communication.

VI. CONCLUSION

In this paper, we designed a light-weight and effective secure access control scheme for IMDs during emergencies. Our scheme utilizes patient's biometric information to prevent unauthorized access to IMDs. The scheme consists of two levels: level-one employs some basic patient biometric information and it is lightweight; level-two uses patients' iris data to achieve effective authentication. In this research, we also made two contributions in iris recognition area: (1) Based on real iris data, we discovered that there is a special bit set - the Discriminative Bit Set. (2) By experiments on real iris data, we demonstrated that iris recognition can be accomplished by comparing only the Discriminative Bit Set (instead of the entire iris code). This decreases the computation overhead of iris recognition by an average of 58%. The experimental results showed that our IMD access control scheme is very effective and has small overhead (suitable for IMDs). Both the FAR and FRR are close to 0.000%.

ACKNOWLEDGMENT

We would like to thank Libor Masek for his source code [19] that was used to generate iris codes from iris images. We did some modifications on Masek's source code.

This research was supported in part by the US National Science Foundation (NSF) under grants CNS-0963578, CNS-1002974 and CNS-1022552, as well as the US Army Research Office under grant W911NF-08-1-0334.

REFERENCES

[1] Medtronic, Inc., "Implantable pacemaker and defibrillator information: magnets," www.medtronic.com/rhythms/downloads/3215ENp7magnetsonline.pdf.
[2] D. Panescu, "Emerging technologies: wireless communication systems for implantable medical devices," *Engineering in Medicine and Biology Magazine*, vol. 27, pp. 96-101, Mar.-Apr. 2008.
[3] D. Halperin, T. S. Heydt-Benjamin, B. Ransford et al., "Pacemakers and implantable cardiac defibrillators: software radio attacks and zero-power defenses," in *Proc. of SP'08*, pp. 129-142, 2008.
[4] P. Inchingolo, S. Bergamasco, and M. Bon, "Medical data protection with a new generation of hardware authentication tokens," in *Proc. of Mediterranean Conf. on Medical and Biological Engineering and Computing*, 2001.
[5] T. Denning, K. Fu, and T. Kohno, "Absence makes the heart grow fonder: new directions for implantable medical device security," in *Proc. of the 3rd Conf. on Hot topics in security*, pp. 1-7, 2008.
[6] E. Freudenthal, R. Spring, and L. Estevez, "Practical techniques for limiting disclosure of RF-equipped medical devices," in *Proc. of Engineering in Medicine and Biology Workshop*, pp. 82-85, 2007.
[7] USPTO Patent Application 20080044014, "Secure telemetric link," http://www.freshpatents.com.
[8] M. R. Rieback, B. Crispo, and A. S. Tanenbaum, "RFID guardian: a battery-powered mobile device for RFID privacy management," *Proc. of ACISP'05*, pp. 184-194, 2005.
[9] K. B. Rasmussen, C. Castelluccia, T. Heydt-Benjamin, and S. Capkun, "Proximity-based access control for implantable medical devices," in *Proc. of ACM CCS '09*, pp. 410-419, 2009.
[10] X. Hei, X. Du, J. Wu, and F. Hu, "Defending Resource Depletion Attacks on Implantable Medical Devices," in *Proc. of the Globecom 2010*, Miami, Dec. 2010.
[11] D. Halperin et al., "Security and privacy for implantable medical devices," *IEEE Pervasive Computing*, vol. 7, pp. 30-39, 2008.
[12] M. Kücken and A.C. Newell, "Fingerprint Formation," *J. Theoretical Biology*, vol. 235, no. 1, pp.71-83, 2005.
[13] A. K. Jain, J. Feng, and K. Nandakumar, "Fingerprint matching," *M. Computer*, no. 2, pp. 36-44, 2010.
[14] L. Masek, "Recognition of Human Iris Patterns for Biometric Identification," Thesis, 2003.
[15] CASIA iris database, http: //www.cbsr.ia.ac.cn/IrisDatabase
[16] A. Ross, "Iris recognition: the path forward," *M. Computer*, vol.43, no.2, pp. 30-35, 2010.
[17] K. P. Hollingsworth, K. W. Bowyer, and P. J. Flynn, " The Best Bits in an Iris Code," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 31, no. 6, pp. 964-973, 2009.
[18] http://www.patentgenius.com/patent/5674259.html
[19] L. Masek and P. Kovesi, "MATLAB Source Code for a Biometric Identification System Based on Iris Patterns," The School of Computer Science and Software Engineering, The Univ. of Western Australia, 2003.