

technicolor



THE SECURITY NEWSLETTER #20

FALL 2011

THE SECURITY

NEWSLETTER

#20

In this Issue

Editorial	2
Be Our Guest	3
The News	4
AES at 126 Bit Strength?	4
Hacking Femtocells	4
Fingertips Heat	5
DigiNotar Certification Authority: Mission Failed	6
Automatic iPhone Eavesdropping	8
Watermarking 3D Movies	10
Where Will We Be?	14
Technicolor Sponsored Conferences	14

Published Quarterly By
Technicolor Security & Content Protection Laboratories
Part of Technology & Research

Technical Editor: Eric Diehl

Editors: Sharon Ayalde

Contributors: Patrice Auffret
Peter Baum
Gwenael Doerr
Alain Durand
Raphael Gelloz
Olivier Heen
Marc Joye
Yves Maetz
Stéphane Onno

EVP: Gary Donnan

Subscribe to the newsletter:
[securitynewsletter\(at\)technicolor.com](mailto:securitynewsletter(at)technicolor.com)

Report vulnerability:
[security\(at\)technicolor.com](mailto:security(at)technicolor.com)

EDITORIAL

When I started my career many years ago, security was not in the scope of consumers' interaction (except for Pay TV). Today, security is playing an increasing role in our daily life. This issue presents some of the most recent attacks that may affect our life.

People seem to take a fatalist stance to the uninterrupted stream of viruses, spam, and security updates of their computers. Will they keep this placid mood once their connected consumer devices, such as connected TV sets, will be under attack? This new generation of devices will become a target of choice for hackers for several reasons. Firstly, new consumer electronics devices are often not well protected (and sometimes not protected at all). Secondly, it is a new playground for hackers with devices using the hackers' preferred operating system: Linux. Surely, some manufacturers will do a decent job of securing their devices, unfortunately, others will do a poor one.

Our daily life relies more and more on security protocols: e.g., SSL/TLS which secures every https transaction. What would happen if such a dominant protocol would be broken? A new attack, with nickname BEAST, challenges this question (see the news of Alain).

This prevalence of the need for security in our daily life raises many questions. Will consumers care about having a secure device and thus pay a premium for it? How will consumers be able to compare the respective quality of the proposed solutions? What will be the user-friendly mechanisms to update aging security? What will be the business models supporting such crucial updates?

When I started my career many years ago, I hired a young, unknown cryptographer: David Naccache. In those early years, he introduced me to serious security. Since then, he has become a worldwide-recognized expert. Thus, I am honored to have him as the guest of this issue.

E. DIEHL
Technical Editor

THE SECURITY

NEWSLETTER

#20

BE OUR GUEST

David Naccache

David, you are affiliated with the Computer Science Department at ENS, you are Professor at Paris 2 and member of the Scientific Advisory Board for Technicolor.

As an expert in the area of cryptography, what do you think about the recent developments in fully homomorphic encryption?

Fully homomorphic encryption allows one, for example, to make search queries using a search engine (like Google) and find information on the web without revealing the query. It also allows one to buy some content without revealing the actual content that was bought.

Fully homomorphic encryption is still in its infancy. When do you think that such techniques will be found in products?

Today, on a regular PC (Intel Core 2 Duo @ 3GHz), the encryption of a bit requires a few minutes, the decryption requires less than 1 second, and morphism operations require, say, a dozen minutes. Key generation, in turn, requires 45 minutes. These figures are for fully homomorphic encryption. In practice (for example for medical applications), researchers recently showed that quasi-homomorphic encryption can be enough. This means that only a certain number of morphism operations need to be performed.

A dozen minutes for encryption, don't you think that this is deterrent for practical applications?

This has to be compared with the development of cryptography on smart cards. In the eighties, a 512-bit RSA signature required a few minutes. This was improved by a factor of 250,000 by Jean-Jacques Quisquater in the CORSAIR chip by using different tricks: clock, DMA, larger multiplier, algorithmic improvements... If the same speed-up factor occurs, fully homomorphic encryption will become commonplace in, say, ten years.

Another remaining issue is the size of the public key: with the best known implementation, it requires 500 kB (which was already greatly improved from 1TB a couple of months ago).

Furthermore, DARPA founded the PROCEED project with several million dollars for the development of fully homomorphic encryption. To sum up, fully homomorphic encryption will surely impact us in the future similarly to what happened with smart cards. The first commercial applications could emerge within 5 years.

If I am correct, you did some work on the subject. What are your contributions?

I mainly focused on the reduction of the public key. We managed to reduce the key from several terabytes (TB) to 1/2 TB and, in a second step, from 1/2 TB to a few kilobytes. But I am confident that further improvements will be found in the near future. This is a fertile area with many people starting to work on it.

What certainly best characterizes your contributions is your creativity. Do you act similarly in your everyday life?

I frequently change the place of my furniture in my apartment. I install and uninstall software on my computer. I always try new stuff and am attracted by things that are unusual. Some time ago, I read that dogs are wolves that remained cubs: they are continuously playing and see in the human being their parents. I think that researchers are humans who never stop playing. They enjoy discovering and are always fascinated by things that they don't know. I believe that people like Adi Shamir and Ron Rivest do not innovate because they are forced to do so but rather because it is their way of living, they play. In my case, I often go outside of my domain to look elsewhere. For example, I recently posted a report discussing how a program can reverse-engineer itself. In the volume I am editing for Jean-Jacques Quisquater, there is an article on a non-articulated robot that moves itself by displacing its center of gravity like a sea urchin. I also used Principal Component Analysis (PCA) as a strategy to distinguish people. The idea was to measure several parameters of a category of people when playing tetris (e.g., fill rate of the rows, speed, ...), run a PCA and compare it with someone playing tetris. So far, I can distinguish players from non-players, but not for example males from females or scientific versus literary people. I have many such examples. If there is something curious, there are chances that I will be interested in.

THE SECURITY

NEWSLETTER

#20



Which directions do you foresee in cryptographic research?

There is the problem of homomorphic encryption which I already discussed.

Another direction is the use of automated proofs for assessing the security of cryptographic protocols. The main advantage I see is that the machine is endlessly expandable and tireless. We may expect that machines will be able to prove theorems that are beyond the scope of what could be done by a human being.

Yet another interesting research direction is that of leakage resiliency. These models are currently investigated by theoreticians but do not yet correctly model real cryptographic devices. On the other side, practitioners develop countermeasures but cannot prove their resistance. A challenge I see is to connect the two worlds.

Thank you!

D. NACCACHE (Ecole Normale Supérieure, Paris)
Interview by M. JOYE

THE NEWS

AES at 126 Bit Strength?

AES is the encryption standard adopted by the US National Institute of Standards and Technology (NIST) in 2001 and used in many cryptographic products. At the rump session of last Crypto conference (Santa Barbara, August 14–18 2011), an international team of researchers announced the first key-recovery attack against AES. The attack has a computational complexity of $2^{126.1}$ for AES-128. An ideal block-cipher should be such that the best key-recovery attack should be exhaustive search. This result shows that 2 bits are lost in the security of AES. Does it change something in practice or does it endanger cryptographic applications relying on AES? No! The remaining security margin is more than sufficient. Technically, the attack makes use of the biclique concept, a method that was originally developed for the cryptanalysis of hash functions. Details can be found in report <http://eprint.iacr.org/2011/449>

M. JOYE

Hacking Femtocells

In July 2011, the Hacker's choice (THC)¹, a hacker group, disclosed an attack against Vodafone UK femtocell service named "the Sure Signal".

Femtocell is a technology provided by internet providers to enable mobile phone calls from the broadband connection of an Internet box. The femtocell is either a separate device connected to the internet gateway or embedded within it.

This "indoor cell" is primarily used to increase the provider's mobile coverage and might be used by the internet gateway owner and authorized passerby.

The THC group opened the box, soldered a serial access and gained access by finding the weak and fixed root password. Once gaining the root access, it is possible to upload custom software to extract secret information and launch various attacks.²

The main issue is that the femtocell acquires the secret subscriber information (encryption keys and authentication vectors) from the core network. An attacker with root access can retrieve the secret material of a customer connected to the femtocell. This secret information enables an attacker to listen to the victim's call, masquerade as her to access her voice mail, make calls, and send messages at her expense.



¹ "the hacker's choice - THC", n.d., <http://www.thc.org/>.

² "vodafone - THC Wiki", n.d., <http://wiki.thc.org/vodafone>.

THE SECURITY

NEWSLETTER

#20

Normally, the scope of this attack is limited by the fact that the femtocell is configured to accept a limited number of phones (e.g., Up to 32). These numbers have to be registered on the Vodafone's website by the femtocell owner. At boot time, the femtocell retrieves this list of authorized phones from the Vodafone's network. This list is then stored as an XML file in the femtocell's file system. THC could easily modify this file to allow any non-registered phone to make calls.

The security of the femtocell specification is weak by design. It embeds a mini radio network controller working similarly as a macro cell although the femtocell is likely to be in hands of malicious users. A skilled hacker may access the encryption keys that the radio network controller manipulates. In 2009, a 3GPP report³ described numerous attacks including this masquerading attack (Attack #10) and considered this attack as very harmful. It proposed a secure storage and execution environment to mitigate this attack. Recently, attacks on the SFR femtocell presented at Black Hat⁴ highlighted the same weaknesses.

R. GELLOZ, S. ONNO

Fingertips Heat

At the 5th Usenix Workshop On Offensive Technologies, Keaton Mowery, Sarah Meiklejohn, and Stefan Savage, researchers from University of California, San Diego presented⁵ an attack to recover the keys typed onto keypads, like those used in ATM cash dispensers.

The attack uses a thermal camera to measure the heat of the physical keys. It may be performed up to one minute after the user dialed the personal code. The researchers automated the image analysis process, thus opening the possibilities for large-scale attacks. However,



³ tr 33.820 v8.2.0 - Security of H(e)NB, Technical report (3GPP, September 2009), <http://www.quintillion.co.jp/3GPP/Specs/33820-820.pdf>.

⁴ Ravishankar Borgaonkar, Nico Golde, and Kevin Redon, "Femtocells: a Poisonous Needle in the Operator's Hay Stack" (presented at the Black Hat 2011, Las Vegas, USA, 2011), <http://femto.sec.t-labs.tu-berlin.de/bh2011.pdf>.

⁵ Keaton Mowery, Sarah Meiklejohn, and Stefan Savage, "Heat of the moment" (presented at the 5th USENIX Workshop on Offensive Technologies (WOOT'11, San Francisco, USA, 2011), <http://dl.acm.org/citation.cfm?id=2028058>.

it still requires a professional thermal camera that is quite costly and not discreet.

The attack is simple to prevent. Firstly, device provider could use metal keypads instead of plastic ones. Indeed, with such material, the measure does not provide reliable data since the overall keypad acts as a thermal reflector. Secondly, the users could also prevent the attack by applying their hand onto the complete keyboard before leaving the ATM. This would heat all the keys equally and make the measure irrelevant.

Y. MAETZ

The Beauty of the Beast?

BEAST is a new attack against SSL3.0 and TLS1.0 (and previous versions) that was recently presented at the Ekoparty security conference.⁶ BEAST allows the decryption of the SSL/TLS encrypted stream using some predictable patterns of the stream, what allows (dramatically) decreasing the number of trials that are required by the attack. The theoretical attack seems to have been first described in 2006⁷, but BEAST implements it for the first time. BEAST is a JavaScript allowing to inject clear data in the bit stream that is to be encrypted and therefore to perform an adaptive chosen-plaintext attack. The attack needs at least an half an hour session to decrypt a 1,000-character long cookie. The attack does not apparently allow recovering the session key but allows rather decrypting the encrypted stream using some properties of the Cipher-Block Chaining (CBC) encryption mode.

Due to its novelty, we did not have time to study the real strength and the actual prerequisites of BEAST. We will come back in our next edition with an in-depth analysis of the beauty of BEAST.

A. DURAND

⁶ "ekoparty Security Conference 7° edición", n.d., <http://www.ekoparty.org/>.

⁷ Gregory Bard, "A Challenging but Feasible Blockwise-Adaptive Chosen-Plaintext Attack on SSL", 2006, eprint.iacr.org/2006/136.pdf.

THE SECURITY

NEWSLETTER

#20

DIGINOTAR CERTIFICATION AUTHORITY: MISSION FAILED

A hacker has compromised the Dutch DigiNotar certification authority (CA). This resulted in the creation of rogue certificates, granting authentication to evil sites in the name of well-known and trusted domains, like Google.com. For DigiNotar, the breach signed the death of the company because major software providers using Secure Socket Layer (SSL) revoked the CA. Before going to the details of the attack, let us describe the trust model behind SSL.

Return to Basics: the SSL Trust Model

This section explains why the revocation of DigiNotar resulted in new versions for all browsers: Firefox, Chrome, Internet Explorer, Safari, Opera, etc.

The SSL trust model

SSL brings secrecy and authentication between two parties, Alice and Bob. Typically, Alice is a browser and Bob is a web server. The authentication part needs a specific trust model. We describe this trust model in the simplest case when Alice tries to authenticate Bob. Bob owns a key pair made of one public key that everybody knows and one private key that only Bob knows. The public key is used for encryption, whereas the private key is mandatory for decryption.

Trust between Alice and Bob

To be sure that the server is Bob, Alice tries to get a proof that this server knows Bob's private key. If Alice already knows Bob's public

key, then this is not too complicated. She just performs the following sequence (simplified): create a secret message, encrypt it with Bob's public key, send it to the server, and waits for the answer. Only Bob can decrypt the secret message of Alice because only Bob knows the private key for decryption. Then Bob sends the secret message back to Alice (with appropriate protections).

Trust in Bob's public key

So far so good, but there is a first weakness: Alice must know Bob's public key and she must be absolutely sure that this is the public key of Bob. Otherwise, an attacker may substitute his own public key. To build this mandatory trust, "someone" securely ties the public key of Bob with the identity of Bob. This should be done in a way that Alice can easily check. The most usual way is to use a certificate.

Trust in Bob's certificate

Not anyone can make the certificate for Bob: it must be a trusted Certificate Authority (CA). Otherwise, an attacker may forge a certificate. The CA checks the identity of Bob, concatenates this identity with the public key, and signs both together. The result is Bob's certificate. The CA has its own key pair: the private key is used for signing Bob's certificate; Alice uses the public key for checking the signature. To sum up: Alice must know the public key of the CA for checking Bob's certificate.

Trust in the Certification Authority

How can Alice be sure that the public key she knows is the actual public key of a trusted CA? A usual method uses a hierarchy: a second CA signs the certificate of the first CA. Both solutions ultimately lead to the second weakness: how does Alice get the certificates for the initial CA?

Trust in the Browser

The most common answer is: the certificates of the initial CA come with the browser. Indeed, each browser has a list of CA that it trusts. Browsers have dozens of trusted authorities but not necessarily the

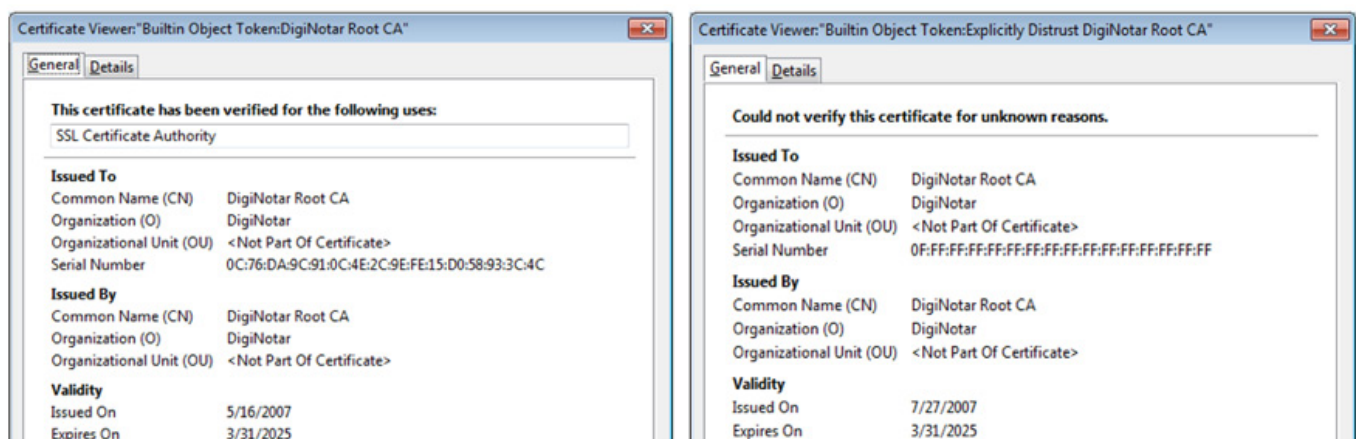


Figure 1: Left: DigiNotar certificate before revocation. Right DigiNotar after revocation still appearing in the list of some browsers (here Firefox) with a serial number set to the maximum, to prevent more complex attacks.

THE SECURITY

NEWSLETTER

#20

same ones. Usually Alice does not choose the CA that she trusts. She just trusts the CA that her browser does trust.

Trust in the place where the browser is downloaded

Eventually, the initial point of trust for Alice is very weak: it is nothing more than the place where she gets her browser. Moreover, when a CA is revoked, as it was the case for DigiNotar, Alice must get an update for her browser or even a new version of her browser depending on how the certificate revocation is managed. In practice, all browser providers updated their browser after the DigiNotar issue.

DigiNotar Intrusion: Timelines

The intrusion was not discovered as soon as it happened. First fail: you shall check your Intrusion Detection System logs (if you have any). DigiNotar reacted only ten days after the issuing of the first rogue certificate, by revoking malicious certificates (the *.google.com was revoked later on). The worst problem was that DigiNotar

In the meantime, major browsers released new versions revoking the DigiNotar root CA. Did it play a role in the death of DigiNotar? Probably. Interestingly, users of Google Chrome were not impacted by this rogue *.google.com certificate. Chrome trusts only a small list of CA as issuers for its domain name. DigiNotar was not part of it.

The Hacker Behind this Intrusion

ComodoHacker claimed this hack.⁹ ComodoHacker pretends to be Iranian with no link to his government.¹⁰ He also pretends to be 21, and to operate alone. Interestingly, he uses the pastebin Web site as a communication vector.¹¹ By reading all his messages, you learn that his first exploit is hacking into Comodo's InstantSSL CA.¹² Also, he pretends to be root on at least four other CAs (he gives GlobalSign and StartCom as examples). His rationale for hacking CAs is that he was not yet able to break RSA, so network intrusion was easier.

The security audit report¹³ shows that known tools were used for intrusion. Furthermore, the hacker developed his own tools (he left

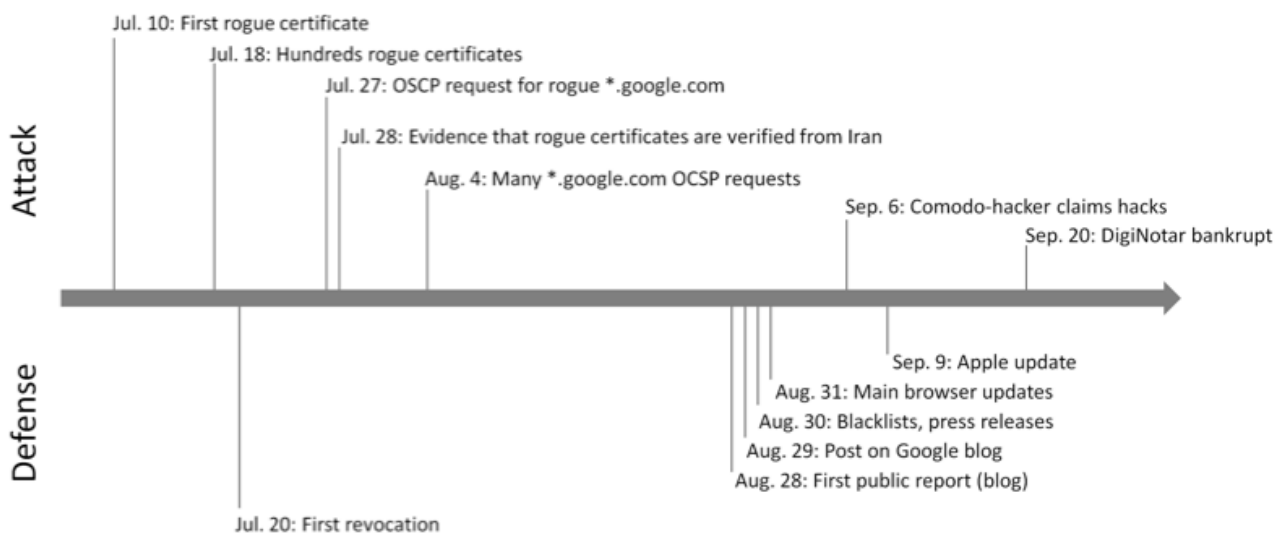


Figure 2: DigiNotar case: timeline

could not find the list of malicious certificates. In fact, some rogue certificates had no records in the CA database of issued certificates. Second fail: you shall have correct disaster recovery procedures. The problem got national when Google found evidence that a *.google.com certificate was issued. Dutch government enrolled a security consulting company⁸ to perform an assessment. After this audit, Vasco closed its subsidiary DigiNotar by filing a voluntary bankruptcy procedure. In certification authority business, thou shall not fail, else thou are dead.

a fingerprint in some of them: Janam Fadaye Rahbar, meaning "I will sacrifice my soul for my leader").

The audit concludes that one unique username/password was enough to control all DigiNotar servers. This account was administrator on a Windows domain with an easy to brute-force password. Frontal Web servers were not using anti-virus that could have easily

8 J. Prins, "DigiNotar Certificate Authority breach 'Operation Black Tulip'" (Fox-It, September 5, 2011), <http://www.rijksoverheid.nl/bestanden/documenten-en-publicaties/rapporten/2011/09/05/diginotar-public-report-version-1/rapport-fox-it-operation-black-tulip-v1-0.pdf>.

9 ComodoHacker, "Striking Back...", PasteBin, September 5, 2011, <http://pastebin.com/1AxH30em>.

10 ComodoHacker, "Another status update message," PasteBin, September 6, 2011, <http://pastebin.com/85WV10EL>.

11 "ComodoHacker's Pastebin", n.d., <http://pastebin.com/u/ComodoHacker>.

12 ComodoHacker, "A message from Comodo Hacker," PasteBin, March 26, 2011, <http://pastebin.com/74KXCaeZ>.

13 Prins, "DigiNotar Certificate Authority breach 'Operation Black Tulip'."

THE SECURITY

NEWSLETTER

#20

detected some of the used tools, and they were not patched with the latest security fixes. Furthermore, an Intrusion Prevention System was in place, but failed to block attacks. Finally, no central logging system was in place. Basic security practices were not in place, and an intrusion was inevitable eventually.

Conclusion

It is all about consequences: one certificate authority is compromised somewhere in the world and then all browsers in the world need to be updated (did you update your smart phone?).

Will similar events happen again? Yes. Because – quoting Bruce Schneier¹⁴– “there are too many single point of trust”. Many certificate authorities are trusted by default without the user even knowing it. In addition, some of these certificate authorities are weak.

Is this the end of the SSL model? Not yet. “SSL is the worst trust model, except for all the others”. In particular, it is widely deployed and easy to use. Other robust models exist, like the PGP model that does not rely on central certificate authorities, but such models are less usable today.¹⁵

P. AUFFRET, O. HEEN

AUTOMATIC iPHONE EAVESDROPPING

Shoulder surfing, i.e., eavesdropping on a victim by looking over his shoulder on his screen and/or keyboard is a well-known, simple, yet efficient attack. With the ever-increasing popularity of smart phones shoulder surfing becomes more valuable for an attacker, since these devices are used for all kind of activities like SMS or email, but also for financial tasks like online banking. Furthermore, these attacks become more practical, since mobile devices are typically used in crowded places where the physical distance between victim and attacker is short enough for an inconspicuous recording of the screen (Figure 3). Paradoxically, the improved user experience, which made the success of smart phones possible, at the same time, simplifies these sorts of attacks. It is another example for the trade-off between usability and security.

Maggi et al^{16,17} describe a shoulder surfing attack on an iPhone, in which keystrokes could fully automatically recognized with a detection rate of up to 97%. They investigated a realistic attack scenario in which, for example in an airport or at a metro station, the attacker can be located directly behind the victim without drawing his attention and records the screen of the victim with a low quality camera like the ones implemented in smart phones. If the victim's phone gives some visual feedback on the pressed key, which is the case for example for iPhone or Android devices, the attacker just needs to know the model of the phone and its keyboard layout to successfully run the attack algorithm.



Figure 3: Attack scenario showing an attacker recording over the shoulder of the victim

The algorithm is divided into three phases. In the first phase, the input video is searched for a possibly distorted keyboard (Figure 4). If it can be found, the picture is registered, i.e. rectified, scaled and cropped, by geometric transformations, so that the result looks as if the picture has been taken by a camera directly above the keyboard (Figure 6). Standard feature extraction and matching methods like SURF (Speeded Up Robust Feature) are used in this step. In the second phase, all non-relevant background is removed from the picture to highlight temporal variation, like moving fingers or the more interesting magnified or highlighted key. In the last phase, the highlighted areas are matched to the known keyboard layout and the best match is selected (Figure 5).

¹⁴ Bruce Schneier, “Forged Google Certificate,” Schneier on Security, September 1, 2011, http://www.schneier.com/blog/archives/2011/09/forged_google_c.html.

¹⁵ Alma Whitten and JD Tygar, “Why Johnny can’t encrypt: A usability evaluation of PGP 5.0,” in Proceedings of the 8th USENIX Security Symposium, 1999, <http://citeseer.ist.psu.edu/whitten99why.html>.

¹⁶ F. Maggi et al., “Don’t touch a word! A practical input eavesdropping attack against mobile touchscreen devices (Milano, Italy: Politecnico di Milano, 2010), http://home.dei.polimi.it/fmaggi/downloads/publications/2010_maggi_volpatto_gasparini_boracchi_zanero_clearshot.pdf.

¹⁷ F. Maggi et al., “Fast, Automatic iPhone Shoulder Surfing”, 2011, http://home.dei.polimi.it/fmaggi/downloads/publications/2011_maggi_volpatto_gasparini_boracchi_zanero_clearshot_poster.pdf.

THE SECURITY

NEWSLETTER

#20




Figure 4: Example of a recorded screen before correction

The detection rate depends on the size and frequency of occlusions for example by victim's fingers and the motion pattern between victim's smart phone and attacker's camera. In the current implemented system, a correct detection rate of up to 97% is achieved. This is similar to humans, who analyze the recorded video manually. However, the automatic key recognition is about 8 times faster, not dependent on typing speed, less tedious and even feasible directly on site.

Countermeasures like switching off the key highlighting or magnification or reducing the viewing angle of the display, possibly combined with an eye-tracking device, decrease the user experience, which is one of the basic success factors of current smart phones. Keyboard



Figure 6: Example of a geometrically corrected picture



Figure 5: Matching and selection of key

phones like BlackBerry are probably less vulnerable, but the market trend seems to go to touch only devices. For touch screen users it only remains to make sure that no camera is pointing at their screen.

P. BAUM

THE SECURITY

NEWSLETTER

#20

WATERMARKING 3D MOVIES

The release of the blockbuster movie *Avatar* in December 2009 is seen as a milestone marking the comeback into vogue of a century-old vintage technology: stereoscopic display. After successive surges in the 50's (anaglyph cinema) and in the 80's (IMAX projection), it is rapidly taking homes by storm with an ever increasing offer of 3D-ready devices e.g., 3D TV sets, 3D set-top boxes, Blu-ray 3D, 3D game consoles, etc. What draws audiences to 3D? "The audience is actually immersed into the world itself. Their feelings are amplified into it, because this is much closer to how we actually see," said Jeffrey Katzenberg, CEO of DreamWorks.¹⁸ Today, the forecast of the 3D market – driven by the demand in sports, games, animation, Sci-Fi, and action movies – has never looked so bright.



A Primer on Stereoscopic Display

Depth perception essentially relies on the fact that the eyes of a human being see slightly different 2D representations of the same 3D scene. When the brain subsequently analyzes those 2D views, it is able to translate the observed differences into depth information. A simple illustrative example is to hold a stretched-up finger in front of you, look at it with one eye open and the other closed, alternating successively between the right eye and the left eye. The closer your finger is from your nose, the more it jerks left and right.

The baseline trick behind 3D movies is simply to mimic this physical phenomenon. At shooting time, the stage is captured by two slightly spaced apart cameras to emulate the typical human inter-ocular distance. The resulting two views (left and right) are then rendered in a way so that each eye only captures a single one of them. Several techniques exist to achieve this:

1. **Color filters.** The right and left views are projected simultaneously using alternate color bands (broadband vs. narrowband) and the viewers are equipped with glasses that exploit color filters to separate the two views.

2. **Polarization.** In this setup, the two views are projected with different light polarization (linear vs. circular) and the glasses of the viewers are equipped with polarized filters to operate the necessary separation.
3. **Active shutters.** The two views are tightly interlaced by projecting successively one frame of the left view and the corresponding frame of the right view. In order to perform the separation, the viewer wears glasses equipped with shutters that alternately obfuscate one of the eyes in synchrony with the projector.
4. **Auto-stereoscopy.** This approach does not require the viewer to wear any glasses. It is mostly used in flat-panel TVs and typically adds an apparatus in front of the display, for instance lenticular lenses or parallax barriers (see Figure 7), to redirect the incoming imagery to several viewing regions at a lower resolution. In other words, only a portion of the displayed pixels is viewable by each eye.




Figure 7: Illustration of the use of parallax barriers and lenticular lenses for auto stereoscopic display (source Wikipedia)

All these solutions are currently in use and offer a different trade-off in terms of cross-talk between the views, loss of luminance, loss of color, freedom of head movement, sweet/dead spots, price, etc.

There are currently two main representation strategies to store and transmit 3D video content between acquisition and rendering times. The first approach simply considers that a 3D movie is a combination of two conventional movies and encodes the two views independently. This strategy is for instance used to produce DCI-compliant 3D movies for digital cinemas. Such a straightforward technique induces however a significant bandwidth overhead. This motivated the design of alternative encoding strategies which exploit the redundancy between the two views in order to achieve a better compression ratio. A typical example is the H.264 multi-view video coding (MVC): both views are encoded using a conventional motion-prediction scheme¹⁹ except that some reference frames, e.g. the intra-frames, of let say the right view are also motion-predicted

¹⁸ Matthieu Aubusson and Vincent Teulade, "Eyes wide open: 3D tipping points loom" (Pricewaterhouse Coopers, October 2009), http://www.pwc.com/gx/en/entertainment-media/pdf/Eye_Wide_Open_3D_Tipping_Points.pdf.

¹⁹ Motion prediction refers to an encoding paradigm routinely used in video compression. Based on the fact that successive frames in a movie look fairly similar, the idea is to predict the current frame to be encoded using already encoded frames. If a block of the frame is in a static area, it can be approximated by the block located at the same position in the previous frame; if a block of the frame is in a dynamic area, it can be approximated by a block at a slightly shifted position in the previous frame.

THE SECURITY

NEWSLETTER

#20

from the corresponding frames in the left view. Alternatively, each couple of frames in a stereo video can be represented as a single view associated to a depth map and an occlusion map. As illustrated in Figure 8, the depth map captures in each pixel the depth of the objects appearing in the view. Equipped with this information and




Figure 8: Image + depth representation. The darker is the depth map, the further away is the corresponding object from the camera.

other parameters (e.g. the distance between the two cameras, the focal distance of the cameras, etc), it is feasible to reconstruct the other view using a simple mathematical relationship. To accommodate for regions visible in the second view that were not present in the reference one, an occlusion map is also appended if necessary.

Tracing 3D Movies

One of the selling points in favor of 3D movies when they started to receive a renewed interest was that they could be the long-awaited silver bullet against piracy. Back in 2009, Jeffrey Katzenberg was indeed claiming, “about 90 per cent of piracy occurs when people bring a camcorder into a screening and shoot it... and that will not work with 3D”. A few years later, the situation looked less ideal than anticipated. Poor quality camcorder copies of 3D projections started appearing on the Internet and proved to be potentially as harmful as their older 2D counterparts.

From a purely business-oriented perspective, any pirate copy that could have a negative impact on box office revenues is considered a threat to the entertainment industry ecosystem. The bad news is that stereo content offers many avenues for generating such pirated samples. Besides the straightforward camcorder, that captures the fusion of the two views, placing one part of the 3D glasses provided by the theater (whatever the underlying technology is) in front of the camcorder lens is a simple trick to isolate a single view of the movie. If the pirate has ‘privileged access’ to the theater screen, she can even record successively the two views. At this stage, she has the choice between (i) combining the two pirated views in an attempt to recreate a 3D movie, or (ii) generating a synthetic view corresponding to a virtual camera placed between the two real cameras that were used when the movie was shot.²⁰ These types of pirate items, 3D or not, can be of good enough quality to be watchable, and are therefore likely to have a notable impact on revenues.

²⁰ Such virtual views can be exploited for instance to adjust 3D rendering to a particular viewing setup (e.g. size of the TV set, distance to the TV set, and so on) and/or the personal taste of the viewer. Virtual views are also used to emulate a virtual camera operated by the viewer when multiple views of the same scene or available e.g. free-viewpoint generation using several views acquired through different cameras scattered across a stadium.

All in all, the situation would not be that bleak if legacy traitor tracing techniques were still effective. In digital cinema, invisible watermarks²¹ are embedded into the video stream at projection time in order to uniquely identify the theater, the theater screen, the date and the time. Should a screening in a digital cinema be camcorderd, a forensic analysis of the pirated sample would allow to pinpoint the cinema owner who let a camcorder enter its premises and to take appropriate remedial actions. At the moment, for 3D movies projected in digital cinema, the same algorithm as the one used for conventional 2D motion pictures is applied to both views. Unfortunately, forensic analysis of real-life 3D pirate samples recently reported that the hid-




Figure 9: Real-life example of a 3D movie using a camcorder without view separation (How to Train your Dragon, Paramount)

den cinema identifier is no longer guaranteed to be retrieved when such a straightforward strategy is in use.

There is no argument about the fact that watermarking 3D movies like a collection of independent videos is really tempting. It allows re-using well tested solutions without having to start from scratch again. However, there are a number of reasons why such an approach was doomed to fail, for the same reasons than watermarking legacy 2D movies does not reduce to watermarking a sequence of independent images.

²¹ Digital watermarking refers to a technology that modifies multimedia content in an imperceptible fashion in order to convey some information in a robust manner. In a traitor tracing scenario, this information is typically the unique identifier of the recipient of the movie. The modifications induced by the watermarking process should not be perceptible by a human being but should be detectable by a machine even if the content has been altered after watermarking. In particular, the hidden information should be recoverable after recompression and display-and-camcord (aka. D/A-A/D conversion).

THE SECURITY

NEWSLETTER

#20

Let us start with perceptual considerations first. It is not because the modifications introduced in each view are invisible when viewed independently that they are also imperceptible when the two views are combined for 3D rendering. From a purely signal processing point of view, watermarking is intrinsically equivalent to adding noise to the video content. Adding the same noise in both the left and right views is equivalent to inserting a static noisy/dirty transparent curtain placed in the middle of the 3D scene with moving objects going through it. In contrast, adding independent noise in the two views is likely to impair depth perception and may result in visual fatigue. Understanding the perceptual impact of noise addition in a particular type of content, in this case stereo video, is a key component of any watermarking system. Such perceptual models are indeed exploited to amplify or attenuate the watermark depending on the region where modifications are introduced. Are changes in static areas with respect to depth less or more perceptible than dynamic ones? Should the background be privileged to the foreground for watermark insertion? Are modifications more perceptible in areas with uniform depth or in regions with steep depth gradient? How do depth perceptual cues interplay with other visual ones e.g. motion sensitivity, spatial frequency sensitivity, contrast masking, luminance masking, etc? Many of these phenomena are not fully understood today and will impact the design of watermarking system for 3D movies in practice.

There is a second reason why it might not be a good idea to readily extend video watermarking algorithm to 3D video and that is robustness i.e., the ability of the system to retrieve the watermark embedded in a pirate sample. We mentioned earlier several types of 3D pirate items (fused views, view isolation, virtual view) and the issue is that each piracy type cast new constraints on the watermarking system. The best case from the watermarking perspective is when the pirate sample is a single isolated view since the full watermarking chain is then equivalent to the one of legacy watermarking systems. Fusing the two views already introduces a new challenge as the two embedded watermarks interfere with each other. Still, best-of-class watermarking systems should be able to 'see' both embedded watermarks in the pirate sample, though with half their original power. The real pain is with virtual views. Under some simplifying assumptions (and neglecting occlusion areas), generating such a virtual view basically comes down to shifting the pixels of a reference view along the horizontal axis by a quantity δ_x referred to as disparity given by the following equation:

$$\delta_x = \frac{f \cdot t_x}{Z} - h,$$

Where f is the focal distance of the cameras (in pixels), t_x the distance between the two cameras (in metric units), Z the depth value of the considered pixel, and h a correcting parameter to adjust the convergence plane. In other words, generating a virtual view induces

a non-rigid geometric transformation of the original view. This kind of distortion is usually not natively supported by legacy watermarking systems and basically disrupts the synchronization between the watermark embedder and its associated detector. It is also common practice to fuse two virtual views generated respectively from the left and right views.

Finally, the watermarking system should not ignore that other applications rely on depth information. Adding the same watermark in the right and left views introduces a bias towards disparity $\delta_x = 0$; inserting independent watermarks in the right and left views introduces noise in uniform disparity areas. If no care is taken, the watermarking system is likely to come into the way of other functions, e.g. depth map estimation, at the risk of potentially disrupting applications that rely on this function.

Emerging Watermarking Strategies for 3D Movies

To accommodate for the previously highlighted challenges, a couple of baseline strategies has been introduced recently to watermark 3D movies.

Strategy 1: View Coherent Watermarking

The first strategy is basically a follow-up on previous works on how to watermark correlated samples. Indeed, in most watermarking systems, there is an implicit assumption that watermarked samples are independent and that they can therefore be watermarked independently. This is typically achieved by applying some transform to the signal prior to the watermark embedding operation. Still, in practice, transform coefficients could still exhibit some inter-dependencies. For instances, successive frames in a conventional video look fairly the same, inducing a high correlation of their transform coefficients. To accommodate for this interdependencies, previous works hinted that the underlying watermark should inherit the self-similarities of the host signal.²²

For 3D movies, this paradigm encourages the watermarking system to guarantee that watermarks embedded in the left and right views are coherent with the associated disparity map. In a nutshell, it is equivalent to simulating a couple of cameras filming a scene that has been physically watermarked. Watermark samples are virtually attached to the physical 3D points of a scene and are exported wherever they are projected in a view. A few watermarking systems

22 Gwenaël Doërr and Jean-Luc Dugelay, "Countermeasures for Collusion Attacks Exploiting Host Signal Redundancy," in Digital Watermarking, ed. Mauro Barni et al., vol. 3710 (Berlin, Heidelberg: Springer Berlin Heidelberg, 2005), 216-230, <http://www.springerlink.com/content/3e3xj3wvntfntq/>.




Figure 10: Average luminance value along the rows of a single frame for different views of a video.

adopting this strategy have been recently proposed.^{23,24} Nevertheless, the corresponding detection procedure requires significant side information such as the original views, the parameters used to generate the synthetic view, etc. If not available, these parameters need to be estimated, possibly leading to some instability of the detection procedure with respect to the accuracy of the estimation process. Such constraints/limitations may not be acceptable in practical applications.

Strategy 2: Watermarking in a View-Invariant Domain

When compensating for a particular signal processing primitive starts to be a hurdle, a trick routinely used in watermarking is to identify a domain invariant to this operation and embed the watermark signal into it. As mentioned earlier, under some simplifying assumptions, the displacements of the pixels between two views reduce to a depth-dependent horizontal offset. This observation motivated having a look at the average of the luminance values of the pixels along the rows in a view. As depicted in Figure 10, since pixels are jerking along the horizontal axis, this average remains stable across the views as long as occluded areas remain negligible.

Once this invariance has been isolated, it is rather straightforward for the man of the art to devise a watermarking system that relies on watermarking this invariant representation of the video.²⁵ The strength

of this strategy is that the embedded watermark can be detected blindly without any need for any auxiliary information and regardless of the 3D piracy type (fused views, isolated view, and virtual view). On the other hand, it also raises interesting challenges in terms of exporting the watermark from the invariant domain back to the pixel domain without introducing significant perceptual artifacts.

Concluding Remarks

It is utopian to think that legacy watermarking systems could be readily applicable to 3D movies. This new medium indeed raises brand new challenges for watermarking both in terms of perceptual fidelity and in terms of robustness to new types of piracy. Early works in this area clearly highlighted two principal strategies to tackle these issues: (i) applying correlated watermarks in alternate view in coherence with the underlying disparity maps, vs. (ii) exploiting an embedding domain immune to view synthesis. Research on 3D movies watermarking is only starting and further advances are required to reach the maturity necessary for deployment in practical applications. In particular, current proposals are limited to watermark embedding in the baseband domain and similar techniques operating directly in the corresponding compressed bit stream are still to come.

G. DOËRR

23 E. Halici and A. A. Alatan, "Watermarking for depth-image-based rendering," in 2009 16th IEEE International Conference on Image Processing (ICIP) (presented at the 2009 16th IEEE International Conference on Image Processing (ICIP), IEEE, 2009), 4217-4220.

24 A. Koz, C. Cigla, and A. A. Alatan, "Watermarking of Free-view Video," IEEE Transactions on Image Processing 19, no. 7 (July 2010): 1785-1797.

25 J. Franco-Contreras, S. Baudry, and G. Doërr, "Virtual View Invariant Domain for 3D Video Blind Watermarking", Proceedings of the IEEE International Conference on Image Processing, pp. 2817-2820, 2011

THE SECURITY

NEWSLETTER

#20

WHERE WILL WE BE?

SATIS, Paris, France, November 8-10, 2011

- Paper presentation: Une brève introduction à la protection de contenus, by Eric Diehl

International Conference on Consumer Electronics (ICCE 2012), Las Vegas, USA, January 13-16, 2012

- Paper presentation: Conciliating remote home network access and MAC-address control, by Stéphane Onno, Christoph Neumann, and Olivier Heen

IS&T/SPIE Electronic Imaging – Media Watermarking, Security, and Forensics XIV (ICIP 2012), San Francisco, CA, USA, January 22-26, 2012

- Paper presentation: Simulating large scale acoustic path benchmarking, by Michael Arnold, Peter G. Baum, Manuel Alonso, Ulrich Gries, and Gwenaël Doërr
- Paper presentation: Forensic characterization of pirated cams: Digital cinema vs. celluloid film prints, by Xavier Rolland-Nevrière, Bertrand Chupeau, Gwenaël Doërr, and Laurent Blondé

C&ESAR, Rennes, France, November 28-30, 2011

- Paper presentation: 802.11 Fingerprinting, by Olivier Heen, Christoph Neumann, and Stéphane Onno

TECHNICOLOR SPONSORED CONFERENCES

3rd IEEE Workshop on Information Forensics and Security (WIFS 2011), Foz do Iguaçu, Brazil, November 16-19, 2011

18^{ème} édition des journées SSI (C&ESAR 2011), Rennes, France, November 28-30, 2011

THE SECURITY

NEWSLETTER

#20

NOTES

EXTENSIVE WORLDWIDE PRESENCE



Vancouver
Hollywood
Indianapolis

Palo Alto
New York
Guadalajara

Mexico
Manaus
Paris

Rennes
London
Madrid

Piaseczno
Rome
Bangalore

Beijing
Bangkok
Sydney

TECHNICOLOR WORLDWIDE HEADQUARTERS
1, rue Jeanne d'Arc
92443 Issy-les-Moulineaux France
Tel. : 33(0)1 41 86 50 00 - Fax : 33 (0) 1 41 86 58 59
www.technicolor.com



© Copyright 2011 Technicolor. All rights reserved. All trade names referenced are service marks, trademarks, or registered trademarks of their respective companies.