

PERSONAL RIGHTS MANAGEMENT (PRM)

Enabling Privacy Rights in Digital Online Media Content

Research Report, July 7, 2005

Lothar Fritsch

Johann Wolfgang Goethe - University of Frankfurt am Main

lothar.fritsch@m-lehrstuhl.de

Klaus Kursawe

Katholieke Universiteit Leuven

klaus.kursawe@esat.kuleuven.ac.be

Abstract With ubiquitous use of digital camera devices, especially in mobile phones, privacy is no longer threatened by governments and companies only. The new technology creates a new threat by ordinary people, who now have the means to take and distribute pictures of one's face at no risk and little cost in any situation in public and private spaces. Fast distribution via web based photo albums, on-line communities and web pages expose an individual's private life to the public in unprecedented ways. Social and legal measures are increasingly taken to deal with this problem. In practice however, they lack efficiency, as they are hard to enforce in practice. In this paper, we discuss a supportive infrastructure aiming for the distribution channel; as soon as the picture is publicly available, the exposed individual has a chance to find it and take proper action.

Keywords: Privacy, Voyeurism, Mobile Phones, Cameras, Watermarking

Introduction

Until the 1990s, public distribution of images could only happen in the press, either in print or in electronic broadcast media. To challenge the unauthorized distribution of an individual's image, a media company could be identified and contacted. Furthermore, the media company usually would know who the photographer was.

With the advent of the Internet as a public communication platform, fast and global distribution of images in public with Web pages became common. Scanned photos then were available from a unknown number of private web pages. The availability of digital cameras reduced the cost and shortened the time it took

to put images online. Still, pointing a digital camera at a person will be noticed in many situations.

The early 21st century introduced mobile phones including digital cameras. The camera lens can hardly be recognized, and now everyone who holds a mobile phone in an individual's surroundings could be taking a photo the same time. Now the individual doesn't know whether there are images on the Web, the individual won't see a camera while being photographed or filmed either. This paper deals with the challenge of protecting one's own image and private issues attached to it. With respect to new mobile technologies and distribution channels, we sketch a privacy threat posed by millions of privately owned cameras in mobile phones.

The legal situation is reviewed, and traditional law as well as recent efforts to tackle the issue with new laws or technological solutions is reviewed.

Next is the definition of the threat, where the attacker and attack scenarios are defined. We introduce a protocol based on watermarking and broadcast channels to enable individuals to take notice when photos are produced around them and search for them on the Web.

Finally, we discuss our solution for PRM and draw conclusions towards the feasibility of the technology on mobile phones with particular respect to already existing digital rights management (DRM) technologies.

Table 1 forecasts sales of camera phones to be over 70 million pieces in 2006.

	2002	2003	2004	2005	2006
Embedded camera phones	1.7%	9%	21%	45%	66%
Camera accessories	0.8%	5.0%	3.0%	2.0%	1.5%
Total mobile camera-capable terminal sales to end users (Thousand number of units)	2,626	14,770	24,657	48,159	70,036

Camera-capable mobile terminal shipments to end users by camera type- Western Europe 2002-2006 (Percentage of total mobile terminal market)

Source: Gartner Dataquest (April 2003)

With massive numbers of camera phones out in the public, photos can be taken at any place. News stories about offenders being caught while shooting photos under women's dresses in public are available from the United States, Japan, Great Britain, Malaysia or even Saudi Arabia. Web sites like Voyeurweb.com have been around longer than camera phones exist to even commercially distribute the content. While this intrusive and offensive use of cameras is regarded illegal in many places in the world, other uses seem to create benefits for society – other news stories tell of offenders being identified thanks to camera phone photos taken by by-standers of a crime. Also, the story of Dutch soccer champion Kluivert tells of new uses of public camera infrastructures.

He broke his team's curfew the night before a game, and was photographed by numerous visitors of a night club – which in the end led to Kluivert being expelled from the team (see Barker, 2003)

Considering the favourable uses of camera phones in public, a solution that does detect, but not prevent taking of photos in public places may seem appropriate.

Technologies for mobile media production and distribution

Today, GPRS and other data transmission technologies for mobile phones allow fast distribution of media content. Build-in software for the creation of web based photo albums (Nokia, Fuji, Kodak online photo albums) and MMS phone-to-phone distribution (Nokia Superdistribution, OMA DRM) offers image distribution opportunities even to camera users who are not familiar with internet data transfer protocols and Web page editing.

Examples of legal context

Several countries enacted laws against unauthorized taking of photos with individuals. More countries are debating legislation that is intended to ban camera phones or their use. Some examples are given below.

In **Germany**, a copyright law (“Kunsturhebergesetz”) protects one's own image against unauthorized publication since Bismark's times. Photos can legally be taken without authorization, but their distribution without authorization – even to small audiences – is illegal. Exceptions are photos taken in public places at events where (press) photography usually happens. Also, individuals of “public interest” (e.g. politicians, actors, celebrities) can be photographed and published with limited restriction (see Dix, 2000).

In **Australia**, under the Commonwealth Crimes Act 1914 - Part VIIB, Section 85ZE it is an offence for “a person to knowingly or recklessly use a telecommunications service supplied by a carrier in such a way as would be regarded by reasonable persons being , in all the circumstances, offensive”. In addition, following the widespread introduction of the internet, state laws were changed to address this issue. For example the Crimes Act in Victoria was amended in 1995 to include the offence of ‘Stalking’. This includes telephoning and sending electronic messages with the intention of causing physical or mental harm.

In **Hungary**, “Taking and transmitting recordings without legal or personal accord is unlawful data handling and can lead to civil, or in some cases penal, responsibility”, according to the Hungarian data protection ombudsman Attila Peterfalvi in an 2003 Reuters' article.

Within the **United States**, many state laws, county and city legislation has been enacted to ban the use of camera phones in public places or in certain situations. The laws ban the use of a camera in a particular location.

In **Japan**, offenders – if caught – face fines of 6 months in prison or ¥500.000.

While many countries do have legislation about camera based privacy invasions and the distribution of photos without consent of the photographed individuals, the question of the enforcement remains. The next section reviews current legal and technological efforts.

Current Solutions

The problem of secret photography has been recognized by most of the involved parties, including the manufacturers, politics and private citizens. Some measures have been taken, though with limited effect.

Tougher laws

As mentioned above, many countries have fortified the right on personal pictures, and increased the punishment for the publication of such.

However, this right may be hard to enforce; the photographed individual may never find out about the publication, or at a time where the picture is too widespread to do anything about it. Also, an offender has to be caught on the scene, before the phone digitally transmits the photo away. Even with laws enacted, an individual's only choice would be to arrest the offender instead of waiting for the police to show up. This is not a setting that helps all members of a society with their rights.

Ban phones

One approach is that places especially subjected to illegal photographing - such as public swimming pools, gyms and Saunas - ban the use of cameraphones altogether. Also, many companies have banned cameraphones to counter industrial Espionage, among them DaimlerChrysler. This has led to the situation that even some cellphone producers banned their own devices from their premises, e.g. Samsung and Motorola.

This approach may be a major inhabitation for normal phone users (though banning cellphones altogether in some places is not a bad idea in the first place), and is only suitable for controlled areas with a high risk of secret photographing. Also, the ban has to be enforced somehow, which may not be easy concerning the small size of camera phones.

“Shutter”- noise

Currently, the most common solution to the problem of secret photographing is to add a sufficient loud shutter-noise – whenever a picture is taken, this can be noticed by the environment. This approach has several disadvantages:

- It is often poorly implemented. For example, if a Sony Ericsson T610 is switched into silent mode, this also turns off the shutter noise.
- Given the noise pollution created by cellphones anyhow, this can add to the annoyance of the technology - especially if MMS traffic (and thus the use of phone-cameras) increases the way the industry hopes
- It violates the privacy of the photographer, as everybody, including people not on the picture altogether, immediately learn about him being present with a camera.
- It is mostly ineffective. Not only can the noise get overheard (due to general noise or the environment, e.g. in a Discotheque), it usually does not help the victim. She can shout at the photographer, but in the average situation she will hardly be able to do anything effective.

Given the difficulty to prevent pictures from being taken without dramatically infringing the rights of harmless photographers, our approach targets the distribution channel rather than the creation of the picture. Thus, pictures can be taken without restrictions. However, the subject of the pictures is made aware that some picture has been taken. Furthermore, should the picture appear on the Internet, she has a realistic chance to locate it at an early point in time, when it is still possible to inhibit the distribution by legal means. As an added value, outside of protecting the victims privacy, this technology can also be used to distribute pictures to interested parties.

Enforce Safe Zones by broadcast

Several businesses develop so-called safe harbor technology which is intended to create zones where a broadcast unit tells camera phones that photographing is forbidden there. Two british companies called Sensaura and Iceberg Systems advertise such a technology: *Safe Haven allows the camera functionality of the phone or other electronic devices to be disabled without affecting any other usage of the device. Safe Haven works by transmitting a signal in a localised environment such as a school, swimming pool, office facility or factory, which disables the camera functionality of devices in the nearby environment. Safe Haven enables digital cameras within a variety of electronic devices to be disabled including camera phones, camera PDA's, digital cameras and multi-purpose MP3 players. (www.sensaura.com, Press Release of Sep 11, 2003)*

While this approach empowers property owners to define zones where photo taking is not permitted, it also restricts a user's freedom of taking pictures with consent in the area.

Other problems are the camera phones already sold to the market and the need to implement the receiver technology into all manufacturer's handsets for an effect. Furthermore, to protect individual rights, one needs a portable unit.

This only could guarantee personal rights independent from someone's property protection policy.

The Privacy Tradeoff

In protecting the personal rights of the person involved in our setting, we have to make a tradeoff between two parties, the person being photographed (the individual) and the photographer. We will now state the minimum rights of each party that should be preserved.

Ideally, the individual should have the right to give consent to every picture she plays a major role in; this is the actual right granted by law in the European Union. This right is hard to enforce technologically, however, as it includes judgment on when a picture is a picture of a person, or just a picture of a marketplace that happens to have people on it.

As a minimum, the individual has the right to know she has been photographed, and to have a chance to get an early warning if the picture is being published, allowing her to take appropriate steps in needed.

As long as he does not infringe any personal rights, the photographer should have the right to take pictures without any major obstacles. In this, the protocol should preferably be passive, and not prevent him from taking pictures unless under well defined and measurable circumstances. Furthermore, the photographer has the right to stay anonymous (as long as he does not infringe anybody else's rights).

Finally, the photographer has the right to modify his device; for example, the camera in a PDA should not stop working if the operating system is modified or replaced.

An Infrastructure for Personal Rights Management

Attack Model

We assume that the attacker does not want to spend much resources into breaking the device. Even with a perfect scheme, such an attacker could easily circumvent our entire system by using a traditional camera with a strong zoom optics, or a traditional mini-camera. Thus, protecting against such an adversary is pointless – the problem is not in the professional voyeurs, but in the wide deployment of photographic devices and the ease of secret photographing.

We do assume, however, that the attacker can do simple modifications to the device and the picture, and that the corresponding instructions will eventually be published on the Internet. At the moment, for example, there exist Internet sources that offer modified operating systems for cellphones to turn off the noise generated while taking a picture.

Our protocol leaves a number of attack points where a sufficiently motivated attacker can escape the scheme. This is unavoidable if we want to protect the rights of harmless photographers as well – unless we treat every owner of a cellphone like a criminal, there will always be ways to escape the scheme.

Outside of making this somewhat harder and therefore less attractive to the masses, our protocol also has its merits if combined with legal measures. Though circumventing the protocol may be possible, it does demonstrate that the photographer has “criminal intent”. Thus, it is easier to distinguish a normally harmless person that just couldn’t resist taking a picture in a particular situation from a semiprofessional voyeur with manipulated equipment.

Basic Protocol

Players

There are three major players in our setting.

The **Photographer (Bob)** is the person taking the pictures. Bob uses a **phonecam**, i.e., a cellular phone with a built in camera. Bob has the rights to not be inhibited while taking pictures and has his identity preserved as long as he does not infringe anybody’s rights. Also, Bob has the right to perform “standard” changes to his cell phone, such as updating the operating system.

The **Model (Alice)** is the person that is photographed by the photographer. Her interest is that she has control over pictures taken of her, i.e., if she is the center of the picture, this picture should (ideally) not be taken without her consent. In our protocol, we grant her a lesser right: If a picture taken of her is published, she gets a fair chance to find out early.

The individual uses a **receiver**, which registers the identities of pictures taken in her vicinity. This could be her own cell phone, but also specialized hardware. The receiver can also be in the infrastructure, i.e., it is provided by external parties, e.g., the owner of a Discotheque or even the GSM operators themselves.

Finally, the **search engine** searches the Internet for picture identities and makes them publicly available. This service is not unlike normal Internet Search engines, but with slightly modified rules.

The Protocol



In the first step, Bob chooses to take a picture of unaware Alice. His camera generates a random picture ID, broadcasts it and embeds it as a watermark into the picture.

Alice's receiver picks up the picture ID and stores it for later use.



When Bob puts the picture on the Internet, specialised search engines find it and index it by the watermarked picture ID. Alice sends requests to the search engine with all picture identities that her receiver picked up, and thus locates the picture taken by Bob.

Hardware Implementation

For our protocol to work, we need to establish a connection between the phonecam and the receiver. We assume that no cell phone manufacturer will be willing to add completely new communication technology into the devices to enable a protocol such as the one presented above. Thus, we restrict ourselves to current hardware. This leaves three general ways to communicate between devices:

Infrared

The big advantage of infrared communication in our setting is that it is directed, i.e., the signal can be sent in a way that only devices in the view of the camera receive it. This comes with at some price. The bandwidth of infrared

communication is fairly low, and the distance over which a signal can be transported may be too small. Also, it causes problems on the receiving side: if the receiver is not directed to the camera, it may not get any signal at all.

In the way infrared ports are implemented today, they are fairly easy to block; it is sufficient to glue an object onto the light. Also, jamming the signal with a strong infrared light is fairly easy, which would block all communication.

The first issue can easily be solved by building the receiver into the camera lens. Thus, blocking the communication would disable the ability to take pictures. The second approach is harder to deal with. It may be possible to design a camera that can not take pictures if exposed to a strong infrared signal, but that may not solve the problem (as the jamming signal may be directed) and allow for a denial of service attack, i.e., preventing all cellphone cameras to take pictures at all (that possibility may be wanted though).

Bluetooth

This is essentially the complement of IR: The disadvantage is that a Bluetooth signal is undirected, thus also devices not in the visual scope of the camera get the signal. However, it is very difficult to jam, and the bandwidth is sufficient even for interactive protocols.

An additional disadvantage is that currently, enabling Bluetooth on a phone may pose a security risk. Recent studies show that many Bluetooth phones are open to attacks that may reveal the entire phone memory, i.e., the address book, the calendar etc. This may even be possible if the phone is not discoverable – the mere activation of Bluetooth is sufficient (see for example <http://www.nwfusion.com/news/2004/0211cracksappear.html>)

Thus, unless the security of this technique can be improved, to protect the privacy of her pictures the individual may have to risk a privacy-invasion on her phonebook.

GSM-Network

Finally, by their very nature cellphones are capable of sending and receiving signals on the GSM frequencies.

Thus, the idea is tempting to use that signal to transmit the necessary information. However, in the current specification, the GSM protocol is ill-suited for device-to device communication. Adding this capacity would require major changes in the GSM standard, which is unlikely to happen for the purpose of protecting people from illegal pictures.

It would be possible to use the basestation as an intermediate, i.e., the photographer's device sends a signal to the basestation, which in turn sends a cell-broadcast to all devices in the area.

This creates new problems. For one, one cell may be too big, noticing many devices that don't have anything to do with the picture altogether. Also, phones

at the same location may be locked into another cell or use a different provider.

All of the above

Of course, one can also think of a combination of those techniques; for example, an infrared flash could be used to inform a device that it should now listen to a Bluetooth signal or a GSM cellular broadcast. If implemented properly, this could combine the advantages of all technologies: As the infrared signal only has to carry a binary signal, the low bandwidth and limited range are not problematic anymore. And as receivers that did not see the flash do not listen to the radio signals, they can be configured to not pick up pictures that don't interest them at all.

Attacks on the Hardware

It is to be expected that some users will try to disable the proposed functionality by manipulating their devices. There are three general approaches that already are used to disable the "shutter noise" in current phones:

Implementation errors. Surprisingly many cellphones have implementation errors that allow to circumvent the protective mechanism. For example, some phones turn off the shutter noise if the entire phone is put in silent mode. For our protocol, it is possible that the transmission may be blocked by deactivating Bluetooth or by using it for communicating with another device while the picture is taken. There is little one can propose to counter the problem, and it is not even clear if these errors do not occasionally happen by intent.

Manipulated Hardware. Some users directly manipulate their cellphones hardware, i.e., by building a on/off switch into the speaker. For our protocol, the manipulation could detach the infrared light or the Bluetooth antenna. These attacks require a certain amount of skill and essentially always invalidate the warranty, so they are unlikely to be used by an average attacker. If they are used, there is again little possibility for defense on the software side.

Manipulated Firmware. For some Cell phones, manipulated firmware is available on the Internet. This firmware then turns off the corresponding functionality. This attack is easy to perform by a broad audience. Once the manipulation has been done, essentially everybody is able to replace the firmware. As the possibility for updates and alternative operating systems is a desired one, this problem will remain. However, cell phone manufacturers have recently started to think about other functionality

that a user may not manipulate, e.g., Nokias Superdistribution and Micropayment. Thus, it is foreseeable that this problem will be solved in the near future, e.g. by using a core-operating system that cannot be changed by the owner and building the real operating system on top of this core, or by TCPA/TCG-like technologies.

Software Implementation

A key point of our scheme is to embed watermarks with a picture identity into the pictures. An occasional collision between two picture identities does not cause significant trouble – it merely poses a minor annoyance to a user. Thus, the picture identity does not need to be excessively long. With a k -bit identifier, we need $1.2v2^k$ pictures for the probability of a collision being $\frac{1}{2}$. With an expected 70 million devices sold by 2006, a 40 bit identifier should be sufficient even for high usage of the cameras. Although there are no firm numbers, a embedding a 40 bit watermark into a picture with $640*480$ pixels is quite realistic. For example, the watermark benchmark by Kutter and Petitcolas Kutter and Petitcolas, 1999 performs the tests with 100 bit watermarks on $512*512$, 24 bit colored pictures.

One of the weaknesses of our scheme is that everybody has to be able to extract the watermark information from the picture. This does not inhibit the privacy of the photographer, as the information is a random string without any meaning. However, it does assist the photographer in attacking the watermark, as he can always verify if his modifications destroyed the information.

Limits of the technology Watermarking algorithms are the most critical part in the suggested infrastructure. Watermarking has its failures – the manipulation of digital images can damage or destroy watermarks if enough effort is taken to do so. Some watermarking technologies are robust against strong image operations such as re-scaling an image, others are not. Instead of focusing on each watermarking algorithm's individual failures, we suggest to review the PRM system as a large-scale system like the media industry's DRM efforts. DRM is very likely to be broken or bypassed by skilled individuals, but a high number of consumers lack the knowledge and energy to do so. PRM can be viewed the same – it has the potential to help individuals to detect a high number of privacy violations except for a few skilled ones. The technology might develop to improve watermarking, too.

Search engines The final part of our protocol is a search engine that allows the individual to locate the pictures on the Internet. Today, there are two major distribution channels:

On the World Wide Web, the search engine could work just like ordinary search engines today. All that is necessary is to extract the Watermark from the

pictures and use it as an index. If the Watermark extraction is computationally easy, this can be done within the normal operation of search engines. For copyright protection schemes, commercial web spiders are already available, such as DigiMarcs MarcSpider (see www.digimarc.com/products/imagebridge/MarcSpider/default.asp). Naturally, counter technologies have been developed that hide the pictures from the spider, for example by splitting it into many small pictures or by embedding it using Javascript. This is another point where a sufficiently motivated attacker can circumvent the scheme, which is hard to deal with unless the privacy of the photographer is inhibited.

On Peer-To-Peer networks, searching is somewhat more difficult. However, even now, lists of checksums of various files on peer to peer networks exist, e.g. www.sharereactor.com. A similar technology could be used to centrally collect picture identities, and thus provide efficient searching also on peer-to-peer networks.

Modifications

Stronger Watermarking

To strengthen the watermarking, some technique could be used that allows only selected parties (i.e., the search engines), to extract the watermark from the picture. The advantage is that it would become more difficult to attack the scheme, as the photographer can not easily verify if the watermark has been successfully removed. However, this would give selected parties the exclusive power to use the scheme. This may be unwanted, and raises the question on who selects these parties.

ID of photographer

To strengthen the protocol, the identity of the photographer or his cellphone could be added to the transmitted signal and/or the watermark. This would significantly decrease the risk that the protocol is broken by the photographer, as the individual learns his identity no matter how and if the protocol is distributed. However, this poses a massive privacy problem, as anonymous photographing would become impossible. A possible solution is to encrypt the identity using a randomized encryption scheme. In this case, the identity would only be revealed if there is sufficient evidence that the picture is illegal, e.g., because it was taken inside a public sauna. Still, it remains an open question who is allowed to decrypt the identities and how misuse on this side can be prevented.

Digests instead of Watermarks

An alternative to embedding a watermark into the picture is to broadcast a digest of the picture after it has been taken. This has the advantage that the picture does not need to be modified at all. However, the digest has to be resilient against picture transformations. To our knowledge, no technique of

building a digest of a picture exists so far that would survive simple modifications to the picture. Nonetheless, for existing peer-to-peer trading systems, hash values of files are used to index media data and reate its quality, e.g. at www.sharereactor.com.

Broadcasting the Picture

In addition to the identifier, a strongly compressed version of the picture could be broadcasted as well. This would inform the individual if there is need to take immediate action, e.g. because a specially compromising picture has been taken or because a credit card has been photographed. On the other side, this costs significant bandwidth, and significantly infringes the photographers privacy.

Conclusion

In recent months, cameraphones have also been used in much more malicious ways than „just“ to invade privacy. Several reports have been published of cases where credit card information has been obtained by secretly taking a picture of the card. With today's cameras, a picture with sufficient quality can be taken from about one meter distance to the card.

Control over one's image is hard to enforce today. Using cryptographic technology and legal regulation in the way suggested above can improve an individual's ability to regain control over his image. As suggested by Alexander Dix (see Dix, 2000), privacy officers and data protection activists can draw new possibilities of privacy management online from the exploitation of technology that has been developed for digital rights management.

Initiatives to enact laws that ban the taking of unauthorized photos are of limited effect when they lack a technological support that supports enforcement and prosecution. On the other hand, users and consumers reject technology that presses restrictions on them. Our suggestion of a detection system for private photos being published on the Internet empowers individuals to detect and act upon violations without putting strong restrictions on cameras and photographers.

In our opinion, great advance for individual privacy can be achieved by applying DRM technology for personal rights management. For DRM, in face of its technological uncertainties, a market for watermarking photographs, videos and music has already developed, as illustrated by the vendor DigiMarc. Photo agencies and photographers rely on watermarking technology to counter unskilled attackers. The same approach is feasible for Personal Rights Management.

Future Work

Many new applications of PRM can be imagined. Researchers might encounter filtering approaches where web hosting companies ensure legality of the images posted on their servers by using filtering mechanisms that prevent publishing photographs with no watermarks.

A photographer could choose to put some form of identification into a watermark to enable photo licence selling.

Electronic government applications such as traffic control and ticketing might benefit from adding time and place information into the photos when they are taken.

References

- Barker, Garry (2003). I spy with my little mobile. *The Age*. <http://www.theage.com.au/articles/2003/12/05/1070351787829.html>.
- Dix, Alexander (2000). Das Recht am eigenen Bild – Anachronismus im Zeitalter des Internet? In *Mediale (Selbst-)Darstellung und Datenschutz, Konferenz des LfD NRW*.
- Friedman, Gary L. The trustworthy digital camera: Restoring credibility to the photographic image.
- Kuttern, Martin and Petitcolas, Fabien A.P. (1999). A fair benchmark for image watermarking systems. In *Security and Watermarking of Multimedia Contents*, pages 226–239. SPIE.