WORKING DRAFT

Privacy-DRM: Data Protection by Copy Protection ?

Rainer Böhme[1] and Andreas Pfitzmann †[2], Technische Universität Dresden, Germany

Abstract

Protecting media contents and enforcing privacy policies seem to be related problems: mechanisms are needed to restrict data processing outside the copyright holder's, respectively data subject's, domain of control. A systematic analysis of requirements, properties, and experiences with digital rights management (DRM) reveals substantial differences between *Content-DRM* to protect media contents and *Privacy-DRM* to protect personal data. Under realistic adversary assumptions, these differences preclude that even if secure Content-DRM existed, it would still not fulfill all essential requirements for Privacy-DRM.

Introduction

Advances in information and communication technology have led to far-reaching social implications. Technology has rocked the music industry's business models, and it becomes increasingly apparent that it is about to wreck privacy. In both cases, technology has been conceived to cure the problems created by technology. The media industry has invested millions in the research of digital rights management (DRM) to restrict uncontrolled distribution of media contents. Also privacy-enhancing technology by ex ante *data avoidance* is reasonably well understood (e.g., [Ada06]), but proposals to *enforce data protection policies* ex post by technical means are less researched. This article revisits the idea of recurring to digital rights management to safeguard personal data online and thereby protect people's informational privacy even in situations where privacy-enhancing technology based on the principle of data avoidance meets its limit.

The article is organized as follows: Section 1 defines the problem and discusses similarities between the application of DRM for the protection of media contents and of privacy, respectively. Successes and caveats drawn from the experience with existing DRM are reviewed in Section 2, before Section 3 evaluates them against a structured account of the

[1] Corresponding author: rainer.boehme@tu-dresden.de
[2] Andreas Pfitzmann died on 23 September 2010. This manuscript will remain a mimeo in its state of late 2009.

specific requirements for privacy protection. As a result of this analysis, Section 4 concludes with a reserved view on the prospects of digital rights management to defeat realistic privacy threats.

1. The vision of Privacy-DRM

Digital rights management systems subsume technologies that enable copyright holders to control the consumption and distribution of their digital contents. In this context, a vision of "self-defending" media data has emerged: legitimate users would be able to play back media on their own devices easily, but the media would turn out unusable in response to attempts of illegitimate use. To realize this end, typical DRM systems consist of components that protect specific stages in the exploitation chain of digital contents. Depending on the architecture, various technologies come into play: typically symmetric or asymmetric encryption, digital signatures, robust and fragile digital watermarks, fingerprinting codes, rights description languages (e.g., XrML, ODLR, NEMO), as well as trusted computing technologies based on tamper-resistant hardware.

In the following, we describe the analogy between well known *Content-DRM* (C-DRM) established for copyright protection and the less known concept of *Privacy-DRM* (P-DRM) to protect personal data [KK03].

The "DRM-problem" for media contents can be stated as follows: content provider *A* wants to make accessible content *C* to client *B* in a specific way, but prevent him from doing *everything* with it. This is difficult, as content *C* leaves the trusted domain of *A* and enters *B*'s domain of control, e.g., by storing the media file on *B*'s computer. The DRM-problem can only be solved if *A* manages to establish a protected area within *B*'s domain of control. The protected areas must be trustworthy for *A* and defeat illegitimate access to *C* (see Fig. 1).

[Fig. 1 about here]

There exists a similar problem in P-DRM: client *B* wants to make accessible personal data *D* to data controller *A* in a specific way, but prevent him from doing *everything* with it. Again, making accessible implies that *D* is processed in *A*'s system, in which now *B* has to establish a trustworthy protected area (see Fig. 2).

[Fig. 2 about here]

Unlike conventional security technology, P-DRM tries not only to protect against illegitimate access by third parties, but also includes the operator of the data center as possible adversary and tries to limit his capabilities. This makes P-DRM a particularly interesting and desirable technique for informational privacy protection in distributed systems. Note that this article does not discuss scenarios that require weaker security guarantees, which might be enforced with standard techniques.

To build a common terminology for both C-DRM and P-DRM, we call the party who "owns" a data object *source* (*A* in the case of C-DRM and *B* for P-DRM) as opposed to *recipient*, which denotes the party in whose domain of control a data object is placed. The source defines rules in so-called *policies* that describe how and under which conditions media contents *C* or personal data *D* may be processed (e.g., played back). These policies are attached to the data object either by embedding or appending. Typical policies for C-DRM bind playback of media contents to specific devices or users, or limit the maximum number of playbacks. For P-DRM, conceivable policies include purpose binding, automatic deletion after elapse of a defined retention period ("digital forgetfulness" [BJ02, May07]), control of transmission to third parties, or regulation of permissible data mining (e.g., disallow specific privacy-invasive techniques, but explicitly allow the compilation of sufficiently anonymized statistics [AP08]).

The responsibility for evaluation and enforcement of policies in specific instances is with the implementation of the protected area within the recipient's system. Complicated policies or those that depend on external conditions, which cannot be determined reliably on the recipient's system, can be decided "online". This requires a communication channel over which permission can be requested from the source prior to each processing of the data object. It is obvious that the integrity of policies and of the communication channel, if applicable, must be ensured. Otherwise the DRM protection could be circumvented easily.

[Fig. 3 about here]

Figure 3 depicts the execution layers of typical multi-purpose computers. Unfortunately, in particular on freely programmable computers it is impossible to effectively protect data objects from access by lower layers. This renders it infeasible to establish secure protected areas for contents *C* and personal data *D* on such devices. Many practical DRM systems therefore realize only imperfect and – e.g., due to "security by obscurity" – temporary protection.

2. Lessons learned from Content-DRM

For more than fifteen years, content providers have been trying to implement DRM systems in practice. Thereby two relevant architectures to design secure DRM systems have emerged. To emphasize the differences, we will briefly review both in their idealized form.[3]

In the early days, DRM pioneers envisaged an *open architecture with traceability*. The main idea of this architecture is to embed source identifiers into media data, for instance by using *robust digital watermarks* [Cox07]. These technologies can mark media contents by altering its signal representation imperceptibly so that a *watermark detector*, parameterized with the respective secret key, can detect the presence of embedded digital watermarks with sufficient certainty. According to common security definitions, it should be impossible for entities who do not posses the secret key to remove or destroy the watermark without, at the same time, degrading the media's quality severely. Robust digital watermarks that convey recipient identifiers are known as *fingerprints*. Creating secure fingerprints is at least as difficult as robust digital watermarking for three reasons. First, fingerprinting relies on watermarking and therefore all attacks to render a robust watermark unreadable must be defeated. Second, fingerprint identifiers must withstand attempts of manipulation, i.e., leaving detectable, but false identifiers in the media signal. Otherwise there is a risk of framing innocuous recipients as suspect traitors. Third, the adversary model has to anticipate that multiple recipients collude and exploit the differences between their individual versions of the media object as side information to locate and eventually invalidate the fingerprint. Hence we can put down that the sole mechanism of open DRM architectures to effectively protect against illegitimate use – in particular uncontrolled distribution – relies on deterrence. Making recipients accountable for illegal copies of media objects marked with their identifier creates strong

[3] Other solutions are conceivable as well and have been tried in practice. Many of them are based on too optimistic assumptions (e.g., no reverse engineering of software) and hence cannot offer protection against strong adversaries. We do not recommend using such approaches as models for privacy protection mechanisms.

incentives to change behavior towards a more responsible use of such protected media contents. For the reasons given above, robust digital watermarking is indispensable to entangle recipient identifiers and media contents irreversibly.

The *closed architecture*, by contrast, realizes security properties by interventions in the execution layers (Fig. 3). One option is dedicated hardware for playback, for instance DVD players or pay-TV decoders. Such hardware is not freely programmable and – provided that it contains secrets of the content providers – must be equipped with physical protection against tampering and reverse engineering. In this case, media objects can be transmitted in encrypted containers without trusting networks nor client computers. Their content is decrypted and played back not before it reaches the dedicated hardware, which forms a trusted domain of the content provider. So-called *trusted computing* [Fel03] techniques can help to establish such trusted domains also on freely programmable hardware, which should be protected against physical manipulations by its owner, though. The most widely known solution, promoted by the Trusted Computing Group[4], and industry partnership, employs a so-called *trusted platform module* (TPM). The module combines functions to store secrets securely – also with respect to the operator of the hardware – and to calculate and manage checksums over programs that reside in the computer's memory. If and only if these checksums match with reference values defined by the content provider, the TPM shall allow the decryption of media objects (usually by employing its secrets, otherwise third parties could decrypt as well). The actual decryption and processing can be implemented in software, since content providers would only allow checksums of programs in which they trust that the data processing does not violate their interests (see Fig. 4). Hence, closed DRM architectures can function without watermarking, but depend both on the security of tamper-resistant hardware (in our example: TPM *and* the surrounding hardware) and the correctness of software.

[Fig. 4 about here]

In practice, none of the two architectures has succeeded in its pristine form. Virtually all commercially viable C-DRM systems combine elements of both approaches [Katz04].

- Unsolved issues of the *closed* architecture include high costs of sufficiently secure special-purpose hardware, the so far hesitant adoption of trusted computing in consumer computers, and the so-called *analog hole*. The latter refers to the possibility

[4] https://www.trustedcomputinggroup.org/home (last access: Oct 2009)

of recipients to record media contents during playback "in place of a sense organ" and, if necessary, subsequent re-digitization.

- Weaknesses of the *open* architecture mainly emerge from the fact that technologies for sufficiently secure robust digital watermarks are not available at all. Supposedly secure watermarking algorithms have been exposed to public scrutiny in a string of challenges, and we are not aware of a single case in which the tested algorithms could withstand the attacks. This is even more noteworthy if one considers that those challenges were conducted under unrealistically difficult conditions (for the attacker). The organizers of the SDMI challenge published less information about the system than a realistic attacker would have at his disposal. The more recent BOWS[5] contests excluded entire classes of particularly strong geometric attacks [PAK98] by choosing an unrealistic quality metric (that is way too sensitive to geometric deformations compared to human perception). Further impediments to the use of fingerprints are that extracted identifiers of recipient [PS96] or source [CM+98] often do not qualify as evidence in court. This thwarts the effectiveness of the deterrence mechanism. Finally, a technology by which the inadvertent loss of a (portable) playback device may trigger claims of indemnification and liability for the redistribution of all media objects marked with one's identifier quite understandably faces acceptance problems.

In response to these problems, additional mechanisms have been conceived to provide supplemental protection:

- Manufacturers of consumer electronics equip their recording devices (video cameras, scanners) with watermark detectors and program devices to deny recording or reproduction of media that contain specific protection marks. For example, this mechanism is known to be in use to impede the digitization of banknotes [Mur04].
- Content providers search or let search the Internet and file-sharing communities for illegal copies of their contents. They try to stop further distribution by means of filters (blocking) or issuing takedown notices to the hosting providers [Cla00].
- The authors of standards for videodiscs have tried to separate markets and limit global distribution of pirated material by using *region codes* assigned by continent.
- Last but not least, the U.S. *Digital Millennium Copyright Act* (DMCA) and comparable legislation in other countries created legal provisions that help content providers to prosecute copyright infringement. It further criminalizes the development

[5] Break Our Watermarking System; see http://bows2.gipsa-lab.inpg.fr/ (last access: Oct 2009)

and distribution of tools to circumvent technical protection measures including DRM. Unfortunately we witnessed cases where the DMCA served as a legal basis for attempts to muzzle academic arguments criticizing the effectiveness of DRM [Cr+01]. Moreover, the definition of what constitutes an illegal "circumvention tool" has turned out to be highly problematic.

In spite of this array of protection techniques, DRM for media contents is not an outright success story. Virtually every protection mechanism has been broken. Owing to the BORE principle (*break once, run everywhere* [And03]), the thus "freed" media contents became available for free at least to computer-literate and patient users with broadband Internet connection. Honestly paying users complain about incompatibilities due to DRM, and became victims of security vulnerabilities (e.g., the Sony rootkit [HF06]). Economists point to the social cost of incompatibilities between competing DRM systems, which hinder competition and lock consumers into single platforms [Eco06]. As a result, the music industry already started rethinking. Apple, the market leader for online distribution of music, has offered DRM-free content for download in iTunes since mid-2007. However, it seems that Hollywood has not fully lost its hopes in DRM, yet [JL07].

3. Personal data is not media contents

Given the similarities between the DRM problems for media contents and personal data, using DRM technologies for privacy protection sounds compelling [KK03], at least on a superficial level. However, a more thorough analysis reveals differences in the requirements between the two applications, which call for distinct technical and organizational measures. Our discussion of differences between C-DRM and P-DRM is structured into (1) properties of the data to be protected, (2) aspects of organization and control, and (3) economic aspects.

3.1 Properties of the data to be protected

Digital media contents differ from digital representations of personal data in at least three relevant respects:

1. share of irrelevance,
2. valuation over time, and
3. sensitivity per bit.

As to point 1: multimedia data, even after state-of-the-art lossy compression, contain a large share of irrelevance, which cannot be clearly separated from the actual information[6] (otherwise this separation could help to improve lossy compression algorithms). Digital watermarking techniques, which are necessary to construct DRM systems according to the open architecture, exploit this imperfect separation and embed encrypted information into those portions of media data that are irrelevant with very high probability. Robust digital watermarks are embedded with redundancy by using error correction codes. As a result, even very short encrypted payloads (e.g., identifiers of source or of recipient) lead to quite a number of changes in the media data. By contrast, personal data (address, bank account details, marital status etc.) are mostly represented in very discrete attributes, i.e., without irrelevance. Therefore it is virtually impossible to embed digital watermarks in them without changing their semantics. Even adding artificial irrelevance to the data yields no improvement since personal data can often be transformed to a canonical form which does not include irrelevance. This would allow separating the information from the watermark (and thus violate a security requirement for robust watermarks). To make an example, consider so-called "address fingerprinting" by means of fictitious middle initials. To find out who shares personal data of John Doe, he could fib a bit and add vendor identifiers to his address: John *A.* Doe for Amazon.com, John *B.* Doe for his bank, and so forth. However, in many countries there exist public registers against which addresses can be matched so that the middle initial would be corrected or disappear – and with it the "fingerprint". Hence, we do not see a viable way to fingerprint personal data except for two specific cases, which we sketch for the sake of completeness: first, the increasing popularity of *biometrics* leads to processing of new classes of personal data, which contain a higher share of irrelevance than discrete attributes. Second, fingerprinting of *collections of personal data* can work. The mark is hidden in the fact whether, e.g., an existing address belongs to the collection or not. List brokers in direct marketing make use of this method to protect their lists against illegitimate (re-)use. The leaser of an address collection does not know which address the list broker has inserted as control entry. However, these special cases to not invalidate our stance that digital watermarking of personal data is impractical in general. So we conclude that P-DRM systems by principle have to be constructed as closed architecture [SS01, IS05].

As to point 2 (valuation over time): the commercial value of media contents tends to decline over time. More and more people have seen a movie, fashion and flavors change, and

[6] Only in this paragraph, the terms information, redundancy, and irrelevance refer to their strict information-theoretic meaning [CT06].

secondary distribution channels are standard, such as samplers for music or television for blockbusters. Many C-DRM systems are designed against the backdrop of this devaluation. For example, the revocation of keys for compromised devices (or manufacturers) only becomes effective for new releases [JL07]. So protection is only enforceable for new contents (bar few exceptions, i.e., if policies are decided online). For personal data, it is very difficult to make reliable predictions on their valuation over time [BB09]. If we consider that humans tend to change their personality only slowly, it becomes apparent that many people may find it harder to deal with personal data of events that lie back far in the past (e.g., foolishnesses of high-school days). Moreover, theories suggesting that automatic "forgetting" would be socially beneficial by partly re-establishing the social reality before the age of digitization also imply that older personal data in the wrong hands is deemed more harmful than more recent information [BJ02, May07]. All this suggests that we should not build P-DRM systems on the assumption that personal data decreases in value like media contents.

As to point 3 (sensitivity per bit): personal data are, relative to their size, much more sensitive than media data. The monetary loss of illegal copies of two audio CDs is about US$ 40, assuming the gross sales price. Even if, say, ten further persons who otherwise would have bought the CD (!) make subsequent copies, the total loss sums up to about US$ 440. In November 2007, the British HM Revenue & Customs authority had to acknowledge that it lost two CDs with personal data including bank account details of all recipients of child benefits in Britain – altogether more than 15 million records;[7] and further events followed in a series of data leaks. According to an estimate of TrendMicro,[8] an IT security consultancy, the black market price for valid combinations of address, bank account, and date of birth was between US$ 80 and US$ 300 at about the same time. Even a very conservative projection based on the assumption that only one in ten records has a value of US$ 80, the total loss reaches the immense amount of US$ 120 million! The by orders of magnitude higher value of personal data makes them targets of more powerful attacks and thus calls for much higher security standards. This means in particular that the *analog hole* must not be ignored for processes that allow the output of personal data. Only few data handling processes require neither any output nor manual intervention. Examples include the collection of aggregated statistics (census surveys), or data escrow, say, for emergency use under unambiguously

[7] See, for example, http://www.lightbluetouchpaper.org/2007/11/20/government-security-failure/ (last access: Oct 2009).
[8] See http://www.informationweek.com/blog/main/archives/2007/02/a_walk_through.html (last access: Oct 2009).

defined and verifiable conditions. Those applications clearly remain exceptions, and for most of them it seems that *data avoidance* techniques offer equally good solutions.

3.2 Organization and control

In a closed DRM architecture, the content provider has to control the protected area on the recipient's system. If the world was as simple as the two-party case discussed here so far, there would be few differences between C-DRM and P-DRM. In reality, however, the situation for C-DRM is as follows. Each of a handful of content providers controls many devices of their clients. If the current practice of incompatible DRM systems continues, then exactly one content provider controls a client device. Alternatively, several if not all content providers could agree to form a coalition, possibly involving hardware manufacturers and vendors of operating systems. This coalition would be built on mutual trust and it can efficiently enforce all participants' common interests. Applied to P-DRM, the first case (incompatibility) is inconceivable due to its prohibitively high cost for all involved parties, and the second case (one or a few broad coalitions) appears quite unrealistic: millions of individuals – partly with heterogeneous interests – would have to form a coalition to control corporate data centers around the world. Nevertheless, proposals that resemble the second case can be found in the literature [KK03, ZD04, SSA06]. One way forward would be to let users delegate the control to a trustworthy data protection authority. This organization would need a mandate and resources to conduct regular audits – also without prior notice – of all data centers that handle personal data. At the same time, one has to ensure that this organization, in fulfilling its critical tasks, remains accountable and does not grow too powerful. Distributing the control functions over several independent organizations could mitigate concentration of power on the audit level. However, it remains an open question if this can be achieved without compromising effectiveness (let alone efficiency). We are not aware of any fundamental research on distributed audits with regard to non-inference constraints.

Also the supplemental protection measures known from C-DRM are only of limited use for P-DRM. Realistically speaking, neither individual users nor a control organization will be capable of actively searching for cases of illegal data sharing in closed business-to-business information systems, or even block data sharing with technical means. Merely the use of region codes appears workable, although their protection is notoriously weak. Region codes at least could technically prevent *honest* data controllers from transferring data *inadvertently* to jurisdictions in which the agreed data protection policies might not be followed.

Lastly and for obvious reasons, new DMCA-like legal provisions to outlaw development or possession of data mining software are not an option. Like the recently introduced ban of so-called "hacker tools" in many countries, this might cause similar uncertainty (How to draw a line between "hacker tools" and dual-use security tools?) and chilling effects on research and innovation.[9]

3.3 Economic aspects

Last but not least, we sketch an economic argument: C-DRM systems are commercially successful if they reduce illegitimate use of media contents substantially; say, reduction by 90% would be a great success. Content providers act rationally if they do not care what the remaining 10% of users do with their content (as long as they do not initiate super-distributions). Let us apply this to privacy: a P-DRM system which makes 90% of organizations and individuals comply with agreed policies is certainly a success compared to the situation today. But it is very questionable if we should call it a great success. In other words: For DRM to improve privacy protection substantially, it must be almost perfect, whereas much weaker security is sufficient to make a substantial contribution to copyright protection. Now consider that the cost of a system grows with the level of security it provides, and with decreasing marginal utility: every additional security investment becomes more costly the higher the level of security already is [GL02]. There are at least two theoretical arguments to back this assumption. First, starting from a completely insecure system, possible security measures vary in their cost. A rational investor implements the inexpensive measures first, so that the more expensive ones remain for later. Second, security measures are typically not independent. The number of interdependencies to be dealt with grows with the number of measures implemented [BM09]. Given this cost function, the content industry will never fund development of DRM technology up to a level of security at which it becomes really interesting for privacy protection. Still, there are doubts on whether the security of DRM can be improved sufficiently at all – even if society would be willing to pay a tremendous price.

4. Conclusion

The similarities between informational privacy protection and protection of intellectual property appear amazing at first sight and almost obviously suggest the adoption of protection technology. In this article, however, a deeper analysis of assumptions and requirements concludes that (1) current DRM technology is very weak, even for the protection of media

[9] See, for example, http://www.lightbluetouchpaper.org/2006/02/10/security-research-may-become-a-crime-in-the-uk/ (last access: Oct 2009)

contents, and (2) Privacy-DRM would have to fulfill even higher requirements under much harder conditions. Robust digital watermarks, a basic technology of current Content-DRM, is practically unusable for Privacy-DRM. So completely different approaches are needed for Privacy-DRM, namely architectures based on tamper-resistant trustworthy hardware in data centers. Even with this, we are not aware of techniques to protect against the threat of *analog holes* for personal data. Since already today, address harvesters copy doorbell panels and let type printed classifieds and telephone books, they will spare no expenses to extract personal data from DRM protected systems and enter them into their own unprotected systems.

The flip side of our requirement analysis is: if Privacy-DRM will ever work securely, then nothing would hinder a complete control of media contents.

Acknowledgements

This article has benefited from constructive discussion with Mike Bergmann, Stefan Berthold, Caspar Bowden, Marit Hansen, Matthias Kirchner, Stefan Köpsell, Jan Schallaböck, and Dagmar Schönfeld. An earlier version of this article has appeared in German in *Datenschutz und Datensicherheit*, volume 35, number 5, 2008.

References

[Ada06] Adams, C. (2006): A classification of privacy techniques. *Univ. of Ottawa Law & Technology Journal* **3**, pp. 35-52.

[And03] Anderson, R. (2003): *Trusted Computing FAQ 1.1*. http://www.cl.cam.ac.uk/~rja14/tcpa-faq.html (last access: Oct 2009).

[AP08] Aggarwal, C. C. and Yu, P. S. (2008): *Privacy-Preserving Data Mining: Models and Algorithms.* Springer-Verlag.

[BB09] Berthold, S. and Böhme, R. (2009): Valuating privacy with option pricing theory. *Workshop on the Economics of Information Security (WEIS)*, University College London. http://www.inf.tu-dresden.de/~rb21/publications/BB2009_PrivacyOptions_WEIS.pdf (last access: Oct 2009).

[BJ02] Blanchette, J.-F. and Johnson, D. G. (2002): Data retention and the panoptic society: The social benefits of forgetfulness. *Information Society* **18** (1), pp. 33-45.

[BM09] Böhme, R. and Moore, T. (2009): The iterated weakest link. *Workshop on the Economics of Information Security (WEIS)*, University College London. http://www1.inf.tu-dresden.de/~rb21/publications/BM2009_IteratedWeakestLink_WEIS.pdf (last access: Oct 2009).

[Cla00] Clayton, R. (2000): *Jury & Judge? How "notice and take down" gives ISPs an unwanted role in applying the Law to the Internet.* http://www.cl.cam.ac.uk/~rnc1/Judge_and_Jury.html (last access: Oct 2009).

[Cox07] Cox, I. et al. (2007): *Digital Watermarking and Steganography*. Second Edition, Morgan Kaufmann.

[Cr+01] Craver, S. et al. (2001): Reading between the lines: Lessons from the SDMI challenge. *Proc. of the 10th USENIX Security Symposium*, USENIX Association, Washington, DC.

[CM+98] Craver, S., Memon, N., Yeo, B. L. and Yeung, M. M. (1998): Resolving rightful ownerships with invisible watermarking techniques: Limitations, attacks and implications. *IEEE Journal on Selected Areas in Communications* **16** (4), pp. 573-586.

[CT06] Cover, T. M. and Thomas, J. A. (2006): *Elements of Information Theory*. Wiley, 2nd edition.

[Eco06] N.N. (2006): Apples are not the only fruit (Economics focus). *The Economist* **380**, 6 July 2006.

[Fel03] Felten, E. W. (2003): Understanding trusted computing. *IEEE Security & Privacy* **1** (3), pp. 60-66.

[GL02] Gordon, L. A. and Loeb, M. P. (2002): The economics of information security investment. *ACM Trans. on Information and System Security* **5** (4), pp. 438-457.

[HF06] Halderman, A. and Felten, E. (2006): Lessons from the Sony CD DRM episode. *Proc. of the 15th USENIX Security Symposium*, USENIX Association, Berkeley, CA.

[IS05] Iliev, A. and Smith, S. W. (2005): Protecting client privacy with trusted computing at the server. *IEEE Security & Privacy* **3** (2), pp. 20-28.

[JL07] Jin, H. and Lotspiech, J. (2007): Renewable traitor tracing: A trace-revoke-trace system for anonymous attack. In J. Biskup and J. Lopez (eds.): *Computer Security – ESORICS 2007.* LNCS 4734, Springer-Verlag, pp. 563-577.

[Katz04] Katzenbeisser, S. (2004): On the integration of watermarks and cryptography. In T. Kalker et al. (ed.): *Proc. of International Workshop on Digital Watermarking*. LNCS 2939, Springer-Verlag, pp. 50-60.

[KK03] Korba, L. and Kenny, S. (2003): Towards meeting the privacy challenge: Adapting DRM. In J. Feigenbaum (ed.): *DRM 2002*. LNCS 2696, Springer-Verlag, pp. 118-136.

[May07] Mayer-Schönberger, V. (2007): Useful void: The art of forgetting in the age of ubiquitous computing. *KSG Faculty Research Working Paper* RWP07-022.

[Mur04] Murdoch, S. (2004)*: Software Detection of Currency*. http://www.cl.cam.ac.uk/~sjm217/projects/currency/ (last access: Oct 2009).

[PAK98] Petitcolas, F., Anderson, R. and Kuhn, M. (1998): Attacks on copyright marking systems. In D. Aucsmith (ed.): *Information Hiding, Second International Workshop.* LNCS 1525, Springer-Verlag, pp. 219-239.

[PS96] Pfitzmann, B. and Schunter, M. (1996): Asymmetric fingerprinting. In U. Maurer (ed.): *Advances in Cryptology - EUROCRYPT.* LNCS 1070. Springer-Verlag, pp. 84-95.

[SS01] Smith, S. W. and Safford, D (2001): Practical server privacy with secure coprocessors. *IBM Systems Journal* **40** (3), pp. 683-695.

[SSA06] Sackmann, S., Strücker, J. and Accorsi, R. (2006): Personalization in privacy-aware highly dynamic systems. *Communications of the ACM* **49** (9), pp. 32–38.

[ZD04] Zwick, D. and Dholakia, N. (2004): Whose identity is it anyway? Consumer representation in the age of database marketing. *Journal of Macromarketing* **24** (1), pp. 31-43.

Figure 1: The “DRM-problem” for media contents

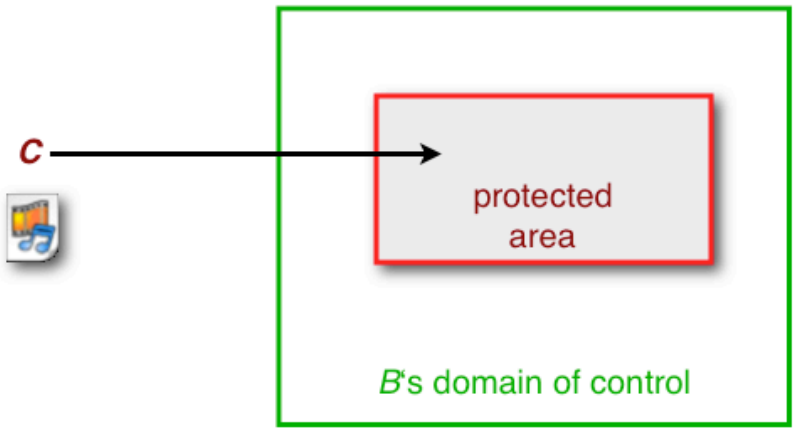

Figure 2: The “DRM-problem” for privacy

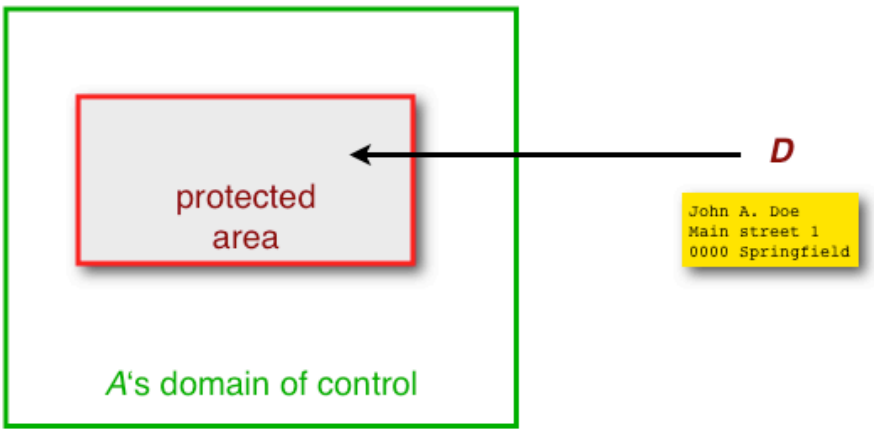

Figure 3: Execution layer: objects cannot be protected effectively from access by lower layers

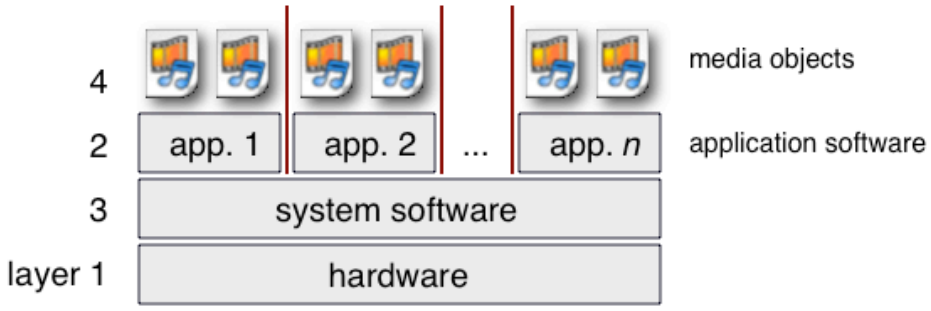

Figure 4: Interaction of a TPM with the execution layers