



Fairness concerns in digital right management models

Luc Bouganim

INRIA Rocquencourt
78041 Le Chesnay, France
Fax: (+33) (0) 1.39.63.55.96
E-mail: Luc.Bouganim@inria.fr

Philippe Pucheral*

INRIA Rocquencourt
78041 Le Chesnay, France

University of Versailles, 78035 Versailles, France

Fax: (+33) (0) 1.39.63.55.96
E-mail: Philippe.Pucheral@inria.fr
*Corresponding author

Abstract: Digital piracy is threatening the global multimedia content industry and blindly applied coercive Digital Right Management (DRM) policies do nothing but legitimise this piracy. This paper presents new software and hardware infrastructure aimed at reconciling the content providers' and consumers' points of view by giving the ability to develop fair business models (*i.e.*, that preserve the interest of both parties). The solution is based on the use of tamper-resistant devices (smart cards) to securely store sensitive data (*e.g.*, personal consumer data or data expressing the terms of a B2C contract or licence) and to perform the computation required by a contract/licence activation. In other words, smart cards can be seen as tamper-resistant Service Level Agreement (SLA) enablers.

Keywords: business ethics; customer loyalty enforcement; automatic and secure B2C contract implementation; Digital Right Management; DRM; XML access control.

Reference to this paper should be made as follows: Bouganim, L. and Pucheral, P. (2007) 'Fairness concerns in digital right management models', *Int. J. Internet and Enterprise Management*, Vol. 5, No. 1, pp.59-77.

Biographical notes: Luc Bouganim is a Researcher at INRIA Rocquencourt. He obtained a PhD from the University of Versailles in 1996 and worked as an Assistant Professor from 1997 to 2002, when he joined INRIA. Bouganim (co-)authored more than 50 conference and journal papers and one international patent and was the co-recipient of five awards. His past research themes were focused on the core of Database Management Systems (DBMS), in particular on query optimisation and execution. Since 2000, Luc Bouganim has been actively engaged in research activities on ubiquitous data management and data confidentiality. He is currently the Vice Head of the Secured and Mobile Information Systems (SMIS) research team. This research team has two objectives: (1) to design embedded database components that can match the

constraints of ultra-light devices, such as smart cards and (2) to devise new architectures that preserve data confidentiality by combining data encryption with security software embedded in secured chips.

Philippe Pucheral is a Full Professor at the University of Versailles and currently in secondment at INRIA Rocquencourt, where he is heading the SMIS research team. He obtained a PhD in Computer Science from the University of Paris 6 in 1991. He (co-)authored more than 50 conference and journal papers, three international patents, four books and was the co-recipient of four awards (EDBT'92, VLDB'2000, e-gate'2004 and SIMagine'05). His domain of interest covers database systems, database components embedded in chips (smart cards), data encryption and smart card-based data protection models (preservation of data confidentiality and privacy, intellectual property and DRM, parental and teacher control). More generally, the research activity of the SMIS team (12 people) focuses on ubiquitous data management and data confidentiality.

1 Introduction

Digital piracy is threatening the global multimedia content industry (there has been a 16% decline in the music market revenues since 1999 (IFPI, 2003)), while forecast institutes agree on the unprecedented potential of the (mobile) audio and video market. Existing Digital Right Management (DRM) models fail, however, in solving this paradox because they badly adapt to several new attractive usage scenarios and because consumers are reluctant to use them for privacy preservation and fairness concerns. Indeed, exasperating coercive methods do nothing but legitimise consumers trying to access digital assets illegally (Champeau, 2004). This paper presents a new software and hardware infrastructure aimed at reconciling the content providers' and consumers' points of view by giving the ability to develop fair business models (*i.e.*, that preserve the interest of both parties).

The solution proposed in this paper capitalises on the democratisation of powerful smart card platforms, which provide an effective element of trust in various client devices (*e.g.*, PC, cell phones, consumer electronics). Initially developed by Bull to secure the French banking system, smart cards are now used successfully around the world in several applications (such as banking, pay-TV, GSM subscriber identification, loyalty, healthcare and insurance). Smart cards have actually reached the highest level of tamper-resistance (EAL-7 security level of common criteria (CommonCriteria, 2005)) and are equipped today with significant computing and storage resources (32-bit CPU, mega-bytes of stable storage) (InspireD, 2006). Hence, these platforms are powerful enough to securely store sensitive data (*e.g.*, personal consumer data or data expressing the terms of a B2C contract) and to perform the computation required by a contract activation. In other words, smart cards can be seen as tamper-resistant Service Level Agreement (SLA) enablers.

The proposed infrastructure, named Mobile Digital Quietude (MobiDiQ), is an XML-based tamper-resistant right-management engine embedded in a smart card. It enforces access control rules (*i.e.*, licences/contracts) depending both on the digital content accessed on the device (music, video, photos, games, *etc.*) and on personal data (historical records, user profile, *etc.*) stored securely on the smart card. Access control

rules based on personal data pave the way for new attractive business models (*e.g.*, they adapt the fee to the exact behaviour of each user). The MobiDiQ access right engine is embedded in the smart card to prevent any tampering from occurring, thereby giving strong anti-piracy guarantees to the content provider. Embedding personal data in the smart card also brings strong guarantees about the user's privacy preservation. In addition, embedding a versatile and powerful access control manager in a smart card gives the opportunity to develop fairer business models. For instance, commercial conditions can be negotiated – then enforced by MobiDiQ – between institutions and content providers to help some categories of citizens (*e.g.*, students) to access valuable contents at a special rate. Parental control rules can also be set up to protect children not only against dangerous contents but also against a prohibited use of legal commercial contents. Privacy-preserving gifting and lending scenarios can be supported as well by MobiDiQ. The expectation is therefore for MobiDiQ to become the mandatory glue between every actor's interest to implement ethical business models.

From a technical point of view, MobiDiQ provides a unique XML-based framework to describe the metadata attached to the protected content, the context of use and the user's profile and to express the access control rules combining them. More precisely, the in-card MobiDiQ engine evaluates DRM access control rules expressed in XPath, the W3C standard language (W3C, 1999), on standard XML descriptions. The support of XML metadata makes MobiDiQ agnostic about the type of multimedia content to be protected and the support of XPath access control rules makes MobiDiQ agnostic about the DRM model to be used at the application level. The latter point is of utmost importance, considering the diversity and absence of interoperability of existing DRM languages and models (*e.g.*, XrML, MPEG-REL, ODRL, XACML, XMCL). Roughly speaking, the MobiDiQ engine can be seen as a DRM virtual machine with XPath access control rules as bytecode. The uniformity of the approach (XML everywhere) greatly simplifies the implementation of the MobiDiQ engine, making it compliant with the current smart card resources in terms of footprint and performance. A prototype of MobiDiQ has been developed on a SIM card platform (cell phone smart card) and has been the recipient of the Gold Award of the SIMagine'2005 international software contest (more than 300 participating teams) (Bouganim *et al.*, 2005).

Thus, the contribution of this paper is twofold. First, it introduces a definition of Fair Digital Right Management (Fair DRM for short) and illustrates it through practical examples. Second, it proposes a tamper-resistant implementation of Fair DRM by embedding an XML access right engine in a smart card.

The rest of this document is organised as follows. Section 2 defines the concept of fair use, explains why commercial DRM attempts have failed so far and highlights the importance of fairness to make DRM models acceptable to consumers. Section 3 gives a global picture of the Fair DRM scenarios MobiDiQ affords. Section 4 introduces the concept of the DRM virtual machine and shows how existing DRM languages can be implemented on top of it. Section 5 presents technical challenges related to the management of XML access rights in a smart card. Section 6 discusses security and performance issues and reports on a preliminary experience with MobiDiQ. Finally, Section 7 concludes.

2 Reconciling fair use and DRM

2.1 The case for Fair DRM

Codified in the 1976 US Copyright Law and frequently used by scholars, journalists and librarians, the fair use provision permits the limited use of copyrighted scientific and artistic material to supplement or briefly illustrate oral or written commentary, literary or artistic criticism, or teaching materials, without permission from the copyright holder. Fair use is necessary to achieve the constitutional purpose of a copyright – to advance knowledge and promote learning. In determining that a use is fair, four factors must be considered:

- 1 the purpose and character of the use (whether it is commercial or non-profit)
- 2 the nature of the copyrighted material
- 3 the amount of the total work used
- 4 the effect of the use upon the potential market (US Department of State, 2006).

To most Europeans, the term ‘fair use’ refers unambiguously to consumer expectations and is not, as in the USA, a legal term of art. In this document, the term ‘fair use’ will be used with a broader meaning, namely a set of good practices participating in the definition of an efficient, competitive and ethical industry of content distribution. In other words, Fair DRM should be the means by which the interest of each party is preserved:

- User’s point of view: as consumers, users are highly concerned about the preservation of the fair use principle as defined in the US Copyright Law. In addition, consumers are expecting to pay for the exact content they are interested in rather than for complete commercial packages. As citizens, users are also concerned about the preservation of privacy (services they use, videos they watch, *etc.*). To illustrate this, a report from IBM-Harris states that the suspicion over the way personal data are exploited by providers is the major obstacle to the development of new applications on the internet (Westin, 1999). As parents, users are more and more concerned about the type of content their children access and the children’s usage of these contents (W3C, 2006). Finally, as members of different communities (family, friends, colleagues, clubs), users expect a reasonable way to exchange (lending, gifting) digital assets.
- Content providers’ and distributors’ point of view: new generation Peer-to-Peer (P2P) technologies accelerate the unrestrained dissemination of content through online networks, regardless of ongoing litigation. The impact of worldwide digital piracy on the music industry (losses estimated at \$5 billion every year by the RIAA) and the movie industry (losses estimated at \$3 billion every year by the MPAA) rose to a dramatic level. According to IFPI (2003), even the sales through the music market declined from \$38.5B in 1999 to less than \$30B in 2003. In 2003, one billion movies and 150 billion music titles were exchanged over the internet. In the absence of a compelling legitimate offer, piracy is becoming the default setting for IP commerce (Tual, 2004). Fair DRM is nothing but another word for ‘compelling legitimate offer’.

2.2 *Fair use and DRM, the impossible marriage*

As stated earlier, digital piracy is threatening the global multimedia content industry in the short term and the production of any cultural assets (even non-lucrative) in the mid term. To face this situation, the major players in the multimedia sector are getting organised and put pressure on their respective governments to enact more coercive laws, like the Digital Millennium Copyright Act (DMCA, 1998) in the USA and the European Union Copyright Directive (EUCD, 2001) in the EU.

However, consumer associations and foundations, both in the USA and the EU, rise up against these laws and the related DRM models imposed by major players. According to Electronic Frontier Foundation (2004), the DMCA anti-circumvention provisions have been used in practice to stifle a wide array of legitimate activities, rather than to stop copyright piracy. DMCA is today accused of chilling free expression and scientific research, jeopardising fair use, and impeding competition and innovation.

Some major distributors tried to integrate fair use practices in their online music service (DRM Watch Staff, 2004). For example, FnacMusic gives its consumers the opportunity to buy either complete albums or single tracks at low charge, raising the bar on consumer friendliness by allowing users to burn each downloaded track up to seven times and transfer it to other devices up to five times. This recalls the war between Microsoft and Adobe over the number of permitted eBook ‘activations’. Music market analysts published a critical analysis of the FnacMusic DRM scheme, arguing that it is both insecure and too coercive, and thus predicting a promising future to P2P piracy (Champeau, 2004).

2.3 *Fair use and DRM, the craved-for marriage*

From the previous discussion, it appears that basic DRM (and simple variants as exemplified by FnacMusic, Adobe or DVD region settings) turns against its initial objective. As mentioned, exasperating coercive methods do nothing but legitimise consumers trying to access multimedia contents illegally (Champeau, 2004).

A Fair DRM scheme is therefore highly required, with the vital objective of reconciling the consumers’, content providers’ and distributors’ interests. A fair DRM scheme must be capable of expressing complex business rules and taking into account contextual information (*e.g.*, user profile, historical data) as a prerequisite to implementing attractive business models while preserving fair use. In the light of the discussion held in Section 2.1, a Fair DRM should therefore exhibit the following properties:

- User friendliness

The consumer should have the opportunity to select the part of a multimedia content he/she is interested in (*e.g.*, a chapter, a music track, a video sequence) and to pay only for that part. Once the content has been legally acquired, the consumer should not be constrained in its daily usage. This means that this content can be made available on any device owned by the consumer, without requiring any complex manipulation (no ‘uninstall from device#1 to reinstall on device#2’). Moreover, the consumer may have the opportunity to freely share legally acquired content with

members of the direct family circle. Lending and gifting must also be possible among small communities of friends, colleagues, *etc.* On the other hand, these practices must strictly prohibit any illegal mass sharing (*e.g.*, P2P sharing).

- Fair superdistribution

Some categories of citizens (*e.g.*, students, needy persons, artists) may have the opportunity to access valuable content at a special rate, assuming commercial agreements took place between institutions (government, universities, associations) and content providers. Again, illegal mass sharing must be strictly prohibited.

- Ethic enforcement

The consumer may have the opportunity to control the contents and his/her children's usage of these contents.

- Privacy

The privacy attached to all the aforementioned practices must be strictly preserved.

The objective of MobiDiQ is precisely to enable such Fair DRM, thanks to the secure smart card platform. The smart card will serve both as a tamper-resistant and private store for contextual information and as a Secure Operating Environment (SOE) to evaluate and enforce complex business rules based on these contextual data.

3 Fair DRM scenarios

3.1 *MobiDiQ infrastructure*

Figure 1 depicts the MobiDiQ infrastructure. Each user has a multimedia cell phone including a MobiDiQ-enabled SIM card. SIM cards are used here as a case study but similar scenarios can be envisioned on any client devices equipped with a smart card or a USB smart token. The SIM card embeds the MobiDiQ access controller and the user's profile (detailed in the next subsections). Basically, the users acquire encrypted multimedia content and encrypted licences from content providers. The licences are decrypted by the SIM card and checked by the MobiDiQ DRM engine, taking into account the metadata describing the multimedia content, the user's profile and potential historical data (all expressed in XML). Depending on the outcome of the licence verification, MobiDiQ may allow 'playing' (part of) the multimedia content. To this end, the decryption key(s) required to decrypt the multimedia content is(are) extracted from the metadata.

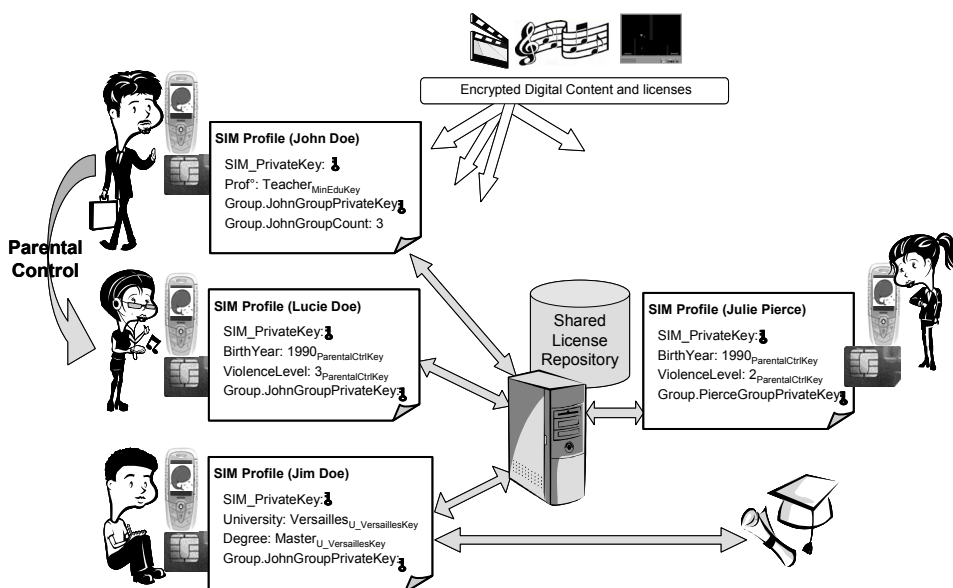
The Shared Licence Repository (SLR) is a cornerstone of the MobiDiQ infrastructure. It ensures four fundamental properties for the user's licences, the profile information and the historical data:

- 1 Availability – the SLR ensures the storage of the aforementioned data and makes them accessible from any user's device (under strict access control policies).
- 2 Resiliency – as a regular database system, the SLR guarantees that the aforementioned data can be recovered in case of failure. This property is required to cope with situations like smart card failure, or loss or theft of the device the smart card is embedded in (*e.g.*, cell phone).

- 3 Sharing – the SLR is a central repository by which licences can be exchanged among a community of users.
- 4 Privacy and integrity – the information stored in the SLR is kept encrypted and signed to guarantee its confidentiality and integrity.

The SLR may be hosted by any server and may be included, for instance, in a telecom operator offer as a (charged or free) service.

Figure 1 MobiDiQ Fair DRM scenarios



3.2 User friendliness scenarios

3.2.1 Private licences

Let us assume that John Doe wants to listen to a given song on his cell phone. A classical way to ensure that John can acquire a private licence is to encrypt it with John's SIM card public key (assuming a PKI infrastructure is used). Therefore, the unique way to play the song is to decrypt the licence with the corresponding private key, owned and protected by John's SIM card. While this solution is satisfactory for private licences used on a single device, it must be extended to enable multi-device usage (*e.g.*, cell phone, mp3 player).

3.2.2 Multi-device usage

To handle multi-device usage, MobiDiQ allows the creation of a group of devices identified by a GroupPrivateKey, and the acquisition of licences encrypted with the corresponding GroupPublicKey. The maximum number of devices belonging to the group, say NbMaxDevice, must be limited to avoid illegal mass sharing. This point will

be discussed further later on. Let us come back to the initial scenario, assuming now that John uses a cell phone and a PDA, both containing a MobiDiQ-enabled SIM card. Thanks to a simple GUI, John asks his phone SIM card to create the group JohnGroup. MobiDiQ then adds two fields to John's profile: (1) the Group.JohnGroupPrivateKey (obtained through a PKI infrastructure) and (2) the Group.JohnGroupCount recording the current number of devices in the group. Adding the PDA to John's group is possible if the condition $(\text{Group.JohnGroupCount} < \text{NbMaxDevice})$ holds. In that case the Group.JohnGroupPrivateKey is transmitted from the phone card to the PDA card (by encrypting it with the PDA card public key) and stored in the PDA card. In addition, Group.JohnGroupCount is incremented.

3.2.3 *Familial usage*

The group mechanism explained above can be used as well to allow media sharing between family members. To prevent illegal mass sharing, it is important to restrict the number of groups a device can belong to. There are two conflicting philosophies here. Existing DRM models do not integrate the group concept, which is nothing but allowing a device to belong to an infinite number of groups. As a side effect, these models drastically reduce the number of permitted transfers of a digital content to any device (*e.g.*, five with FnacMusic), strongly hurting user friendliness. The MobiDiQ philosophy is fairly different. Each device can join a rather reduced number of groups, typically a single one corresponding to the direct family circle (parents and their children).¹ Consequently, the NbMaxDevice limit can be set to a much larger value (*e.g.*, 10 to 15) to encompass all the family's devices, and the number of transfers can be unlimited. Thus, user friendliness and anti-piracy concerns can be met altogether.

3.2.4 *Licence lending*

Binding a licence to a single group of devices can be considered too coercive by many users willing to lend digital content to people outside the direct family circle (*e.g.*, grandparents, friends), as is possible today with classical media (*e.g.*, music CD). MobiDiQ affords lending in such a way that the lending action can remain free of charge (user friendliness) while preventing mass piracy. Let us assume that Lucie wants to lend a given licence to her friend Julie for a week. MobiDiQ handles this case in three steps. First, a specific record is added to Lucie's profile (on her SIM card and on the SLR) to disable a personal use of that licence for one week.² Second, the licence is downloaded on Lucie's SIM card, decrypted, updated to include a one-week validity limit, re-encrypted using Julie's public key, and finally uploaded on the SLR. Julie can now retrieve her private licence and listen freely to the music for one week.

3.3 *Superdistribution scenarios*

Let us now assume that John is a teacher and, as such, may benefit from fee reductions to licences associated with academic contents under an agreement concluded between content distributors and the Ministry of Education. John's profile is enriched with the 'Teacher' tag, authenticated thanks to a cryptographic signature from the Ministry of

Education. In addition, the Ministry of Education delivers to John a specific teacher licence that will be required to validate the reduced fee licences purchased by John. This second licence includes an additional condition checking that John is actually a teacher employed by the Ministry of Education (*i.e.*, checking that the signed ‘Teacher’ tag is present in John’s profile).

The same process may occur with Jim, who is a student at the University of Versailles and benefits from preferential access to some contents (negotiated by the University of Versailles). The university may enrich Jim’s profile with specific information (student status, Master’s degree, education domain, *etc.*) that is used to personalise the licences delivered by the university.

3.4 Ethic enforcement

Finally, MobiDiQ affords protection to children against dangerous contents and also against a prohibitive use of legal commercial contents. Indeed, parental control rules can be defined and can complement regular licences following the same principles as above.

To enable parental control on Lucie’s cell phone, John (Lucie’s father) enters a special PIN code allowing him to define specific fields in Lucie’s profile.³ These fields can be, for instance, Lucie’s birth date (to enforce controls on age limit), some threshold values for violent or sexual scenes in videos, and more generally any kind of information required by John to set up the desired parental control. Then, John selects access control rules that can be based on the media metadata (*e.g.*, video scene description), on contextual data (hours of the day, localisation)⁴ and on historical data (*e.g.*, number of movies or games played on the cell phone), thus allowing powerful and personalised parental control. Parental control looks like a regular licence, the rules of which take precedence (*i.e.*, have a higher priority) over any other rules.

Note that ethic enforcement can also be of interest in the fair superdistribution context (*e.g.*, an association or club negotiates preferential access to some resources for all its members but has strict obligations with respect to its under-18 members or any other category of members).

4 A DRM virtual machine

4.1 About the DRM standards

Several initiatives (*e.g.*, XrML (2006), MPEG-REL (ContentGuard, 2004), ODRL (2006), XACML (2006), XMCL (2001)) demonstrate the need for expressive and extensible DRM languages capable of implementing a large variety of business models. Some of these initiatives are gaining wide acceptance. For example, XrML from ContentGuard is used by Microsoft in its DRM implementations. XrML also formed the basis for MPEG-REL, the Rights Expression Language of MPEG-21. The Open Digital Rights Language (ODRL) has been adopted by the Open Mobile Alliance (2006) for its DRM standard. While these DRM languages have not been designed with fairness

concerns in mind, their expressive power makes them more adapted to express Fair DRM scenarios. Unfortunately, this disparate offer plays against DRM models' interoperability and these standards are too complex to be implemented in smart cards.

4.2 *The DRM virtual machine approach*

While different in their syntax and usages, the DRM languages mentioned above share strong commonalities. To illustrate this, let us consider XrML as a reference language. The constituents of an XrML grant (the central part of an XrML licence) are:

- the principal to whom the grant is issued
- the right that the grant specifies
- the resource that is the direct object of the 'right' verb
- the condition that specifies the terms, conditions and obligations under which the right can be exercised.

Principal, right and resource are respectively named party, right and asset in ODRL and subject, action and resource in XACML, with similar meanings. ODRL integrates conditions within the right statement while XACML distinguishes between conditions and obligations. XACML also supports denials (*i.e.*, negative authorisations) in addition to grants. The way a right is actually exercised is implementation dependent and may differ depending on the DRM infrastructure, on the application and on the type of content to be protected.

By providing a unique XML-based framework in order to describe the metadata attached to the protected content, the context of use and the user's profile, and to express the access control rules combining them, MobiDiQ is a unifying technology. Indeed, the MobiDiQ engine evaluates DRM access control rules expressed in the W3C standard language XPath (W3C, 1999) on XML descriptions. Thus, the DRM languages mentioned above can be easily supported by translating native expressions into XPath, and any kind of content described in regular XML can be protected. This fundamental feature will be exemplified in the next section.

Roughly speaking, the MobiDiQ engine can be seen as a DRM virtual machine with XPath access control rules as bytecode. To help understand how this DRM virtual machine works, we introduce below a brief background on XML and XPath and the XML access control implemented by MobiDiQ.

4.2.1 *XML background*

XML, the Extensible Markup Language promoted by the W3C, has become a *de facto* standard for the presentation, exchange and management of information. Figure 2 presents two samples of XML metadata describing an MPEG-21 video⁵ and a user profile in their textual form. Roughly speaking, an XML document can be seen as a tree of elements (*e.g.*, Seq), each one demarcated by an opening and closing tag (*e.g.*, <Seq> and </Seq>). Attributes (*e.g.*, value) may be attached to elements. Terminal elements (at the leaves of the tree) are represented by text (*e.g.*, Closer).

Figure 2 XML metadata, profile and licences

```

<Video>
  <Title> Closer </Title>
  <Film>
    <Seq>
      <Desc°> ..... </Desc>
      <SexRating> 3 </SexRating>
      <Key> xxxxxxxxx </Key>
    </Seq>
    <Seq>....</Seq>
  </Film>
  <Bonus>
    <Seq>....</Seq>
    ....
  </Bonus>
  <Analysis>
    <Seq>....</Seq>
  </Analysis>
</Video>

```

Video XML Metadata

```

<Profile>
  <SIM_PrivateKey> xdxdd </SIM_PrivateKey>
  <UV_Student> xabc </UV_Student>
  <UV_Master> shdq </UV_Master>
  <Group value = "John"> JohnGrpPrivateKey </Group>
  ....
</Profile>

```

Profile XML Data

Video Licence:

- Require** University_Versailles License
- Rule R1:** < UV_Member, play, ⊕, /Video/Film>
- Rule R2:** < UV_Member, play, ⊕, /Video/Bonus>
- Rule R3:** < UV_Member, play, ⊕, /Video/Analysis>

University_Versailles Licence

- Rule R4:** < [not /Profile/UVStudent], play, ⊕, /Video/Bonus>
- Rule R5:** < [not /Profile/UVMaster], play, ⊕, /Video/Analysis>
- Rule R6:** < ALL, play, ⊕, //Seq[SexRating > 3]>

Licenses

4.2.2 XPath background

Queries can be expressed over an XML document using the XPath language. An XPath expression allows one to navigate in the document through the parent axis (denoted by /) and the descendant axis (denoted by //) and to apply predicates on elements and attributes. The result of an XPath expression is an element (or a group of elements) along with its (their) subtree(s). To illustrate the power and simplicity of XPath, let us consider the following two expressions: /Video/Film/Seq/Key selects all the decryption keys of the sequences of the film, while //Seq[SexRating>3] selects any sequence (anywhere in the document) having a direct child SexRating whose value is greater than 3.

4.2.3 *MobiDiQ access control model*

Several authorisation models have been proposed recently for regulating access to XML documents. The MobiDiQ access control model is inspired by Bertino *et al.* (2001), Damiani *et al.* (2002) and Gabillon and Bruno (2001). The MobiDiQ access control model keeps the foundation and the expressive power of these models while discarding subtleties which could compromise a smart card implementation of the model.

In our model, access control rules take the form of a quadruple <subject[condition], action, sign, resource[condition]>. ‘Subject’ identifies the user(s) the rule applies to. It takes the form of an XPath expression allowing the selection of a group of subjects satisfying given conditions expressed on their profile. ‘Action’ is self-explanatory. ‘Sign’ denotes either a permission (positive rule) or a prohibition (negative rule). ‘Resource’ corresponds to elements or subtrees in an XML document, identified by an XPath expression. We consider here a rather robust subset of XPath denoted by XP{[,*,//} (Miklau and Suci, 2002). This subset, widely used in practice, consists of node tests, the child axis (/), the descendant axis (//), wildcards (*) and predicates or branches [...]. For example, a rule of the form </Subject/Profile[age<16], play, θ, /Video//Seq[SexRating > 3] states that teenagers are prohibited (negative rule) from playing any video sequences having a sex rating higher than 3.

The cascading propagation of rules is implicit in our model, meaning that a rule propagates from an XML element to all its descendants in the XML hierarchy. Owing to this propagation mechanism and to the multiplicity of rules for the same user, a conflict resolution principle is required. Conflicts are resolved using two policies:

- 1 Denial-Takes-Precedence, which states that if two rules of opposite signs apply to the same element, then the negative one prevails
- 2 Most-Specific-Object-Takes-Precedence, which states that a rule which applies directly to an element takes precedence over a propagated rule.

4.3 *Translation of existing DRM languages*

Let us now see how the aforementioned DRM languages can be translated in the MobiDiQ bytecode. While MobiDiQ contains advanced features that cannot be expressed in existing DRM languages, ensuring a backward compatibility is an important concern. We illustrate below the flavour of the translation considering XrML as the target DRM language.

The translation takes place as follows:

- Principal – while a principal must be resolved to a single party (*e.g.*, represented by a private key in a PKI infrastructure) during the interpretation of the right expression, it can actually represent a group of persons (*e.g.*, a member of an institution). In MobiDiQ, the principal (*i.e.*, subject) is expressed as an XPath expression applied to the user’s profile embedded in the smart card. This approach has two main advantages. First, it gives a very simple and powerful way to define principals (*e.g.*, teenagers, Master’s degree students of University U). Second, the smart card guarantees both the privacy of this profile information and its tamper resistance (the card holder could try to tamper her profile to get larger rights).

- **Right** – the right is the permitted action on the resource (*e.g.*, play, print, copy, lend). In MobiDiQ the right (*i.e.*, action) is executed for each XML element qualified by the XPath expression on the resource.
- **Resource** – the resource is the digital asset to be protected. In MobiDiQ, the resource is once more represented by XPath expressions applied to the XML metadata describing the resource. The expressiveness of the XPath language associated with the capability to declare multiple positive and negative rules allows us to define a resource as a set of digital assets (*e.g.*, all movies of category U) or as sub-part(s) of a digital asset (*e.g.*, movie M except the violent scenes) at a very fine granularity level. This fine-granularity level participates in the satisfaction of user friendliness by giving the ability to define very precisely the part of a multimedia content the consumer is ready to pay for.
- **Condition** – conditions can range from simple ones (*e.g.*, a time interval) to rather complex ones (*e.g.*, based on historical data). In MobiDiQ, every element that participates in a condition is integrated in the user's profile embedded in the smart card. Thus, conditions can be expressed in the Subject field again using XPath expressions applied to profile data. It is worth noting that the privacy preservation and tamper resistance provided by the smart card to these data are of utmost interest.

4.4 Running example

Figure 2 shows a sample of the XML metadata attached to a given video, the user's profile and the access control rules expressing the licences downloaded by the user. It illustrates the scenario presented in Section 3.3. The video is divided into several tracks (film, bonus, analysis), each one subdivided into sequences that include descriptions, values indicating the rating in terms of violence, sex content, decryption keys, *etc.* The user's profile is also stored as an XML file with a very simple structure.

In this example, MobiDiQ has to deal with two licences. The first one is issued by the content provider and states that any member of the University of Versailles may have the right to play the Film (R1), Bonus (R2) and Analysis (R3) track of the video. The second licence, delivered by the University of Versailles, adds some restriction to the previous one, specifying that the Bonus track is restricted to students (R4) while the Analysis track is restricted to Master's degree students (R5). Finally, the last rule expresses that any sequence rated with a value higher than three for sex content should not be played. The Require statement stipulates that the second licence is mandatory for enabling the first one. Note that a required licence always restricts the possibilities for the user (Rules R4, R5 and R6 have a negative sign). Indeed, the university cannot grant more rights than those delivered by the content provider itself.

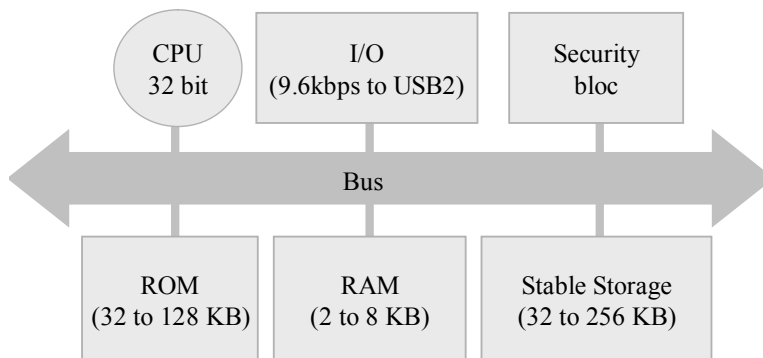
5 Technical challenges

5.1 Smart card architecture

Smart card is today the most widespread representative of an SOE. Existing smart cards typically embed on a single chip, a 32-bit RISC processor (clocked at about 50 MHz), memory modules composed of ROM (about a hundred KB), static RAM (some KB) and electronic stable storage (hundreds of KB of EEPROM or FLASH), and physical security

modules. The limited amount of on-chip computing and storage resources is primarily owing to security reasons: the smaller the silicon die, the most difficult and costly the tampering attacks are. The ROM is used to store the operating system. The RAM is used as working memory (heap and stack). Electronic stable storage is used to store persistent data and downloaded programs. Figure 3 pictures this typical hardware architecture.

Figure 3 Typical smart card architecture



The smart card internal resource balance is very unique. The processing power, calibrated to sustain cryptographic computations and enforce security properties, is oversized with respect to the other resources. Conversely, only a little amount of RAM remains available for the embedded applications, the major part of the RAM being preempted by the operating system. Finally, electronic stable memories generally exhibit very good read performance (*e.g.*, 60–100 ns/word in EEPROM) but suffers from dramatically slow write time (*e.g.*, about 10 ms/per word in EEPROM).

The challenge is then to design a MobiDiQ engine which can accommodate this particular resource balance while providing acceptable performance for the targeted applications.

5.2 *MobiDiQ engine*

The core of the MobiDiQ engine is the XPath access right controller embedded in the smart card. While several access control models for XML have been proposed recently, few papers address the enforcement of these models. Existing solutions rely on a memory materialisation of the input XML document to be protected (Bertino *et al.*, 2000; Damiani *et al.*, 2002; Gabillon and Bruno, 2001). Roughly speaking, these algorithms work as follows. First, a *building phase* parses the input XML document and builds a DOM representation. Second, a *tree labelling phase* evaluates the access control rules related to a given subject and tags any node targeted by the corresponding XPath expression. Third, a *conflict resolution phase* resolves potential conflicts among rules targeting the same node and propagates the final decision about the outcome of each node to its subtree. Finally, the *pruning phase* discards every node annotated negatively.

Considering the smart card constraints in terms of RAM size and write cost in stable memory, building a materialisation of the XML document is precluded. Thus, we designed a streaming XML access right controller. To the best of our knowledge, this work was the first to consider a streaming management of access control policies (Bouganim *et al.*, 2004).

At first glance, streaming access control resembles the well-known problem of XPath processing on streaming documents. There is a large body of work on this latter problem in the context of XML filtering (Diao and Franklin, 2003; Green *et al.*, 2003; Chan *et al.*, 2002). These studies consider a very large number of XPath expressions (typically tens of thousands). The primary goal here is to select the subset of queries matching a given document (the query result is not a concern) and the focus is on indexing and/or combining a large amount of queries. One of the first works addressing the precise evaluation of complex XPath expressions over streaming documents is Peng and Chawathe (2003), which proposes a solution to deliver parts of a document matching a single XPath. While access rules are expressed in XPath, the nature of our problem differs significantly from the preceding ones. Indeed, the rule propagation principle along with its associated conflict resolution policies (see Section 4.2) makes access rules not independent. The interference between rules introduces two new important issues:

- 1 At parsing time the evaluator must be capable of determining the set of rules targeting a given node and deciding which one applies according to the conflict resolution policies.
- 2 Some rules may be inhibited by others according to the conflict resolution policies; thereby optimisations such as suspending evaluations of rules can be devised.

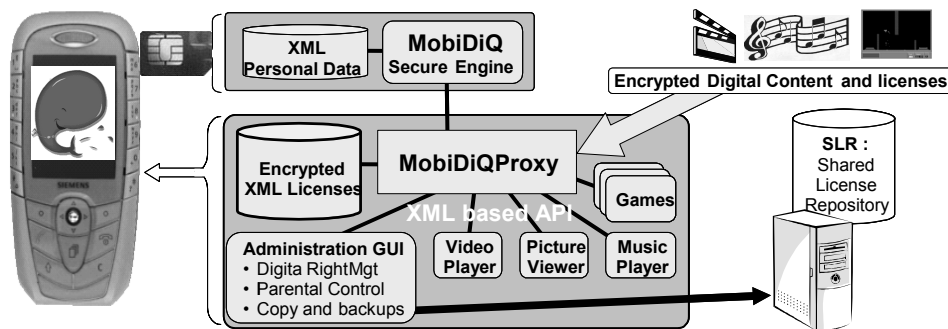
Our streaming access right controller works as follows. The controller is fed by an event-based parser (SAX, 2004) raising open, value and close events respectively for each opening, text and closing tag encountered in the input document. Each XPath expression participating in an access control rule definition (either in the subject or resource statement) is represented by a non-deterministic finite automaton (Hopcroft and Ullman, 1979), named Access Rule Automaton (ARA). An ARA is made up of states connected by transitions. Tokens traverse the ARA while transitions are triggered, at document parsing time. An ARA has one target final state (representing the element targeted by the XPath expression) and may have zero, one or more predicate final states (one for each predicate involved in the XPath expression). When all final states of an ARA have been reached by a token, the corresponding XPath expression becomes active, meaning that it qualifies all forthcoming elements. The outcome of the current resource element is actually determined thanks to a conflict resolution algorithm managing the priorities among all active rules (*i.e.*, access control rules having active XPath expressions).

As a conclusion, the core of the MobiDiQ engine (*i.e.*, the DRM virtual machine) is made up of one XPath evaluator algorithm managing the non-deterministic finite automata and one conflict resolution algorithm. These two algorithms are simple and compact enough to be embedded in existing smart cards. The details of the algorithms can be found in Bouganim *et al.* (2004).

5.3 MobiDiQ architecture

Figure 4 illustrates the components involved in MobiDiQ, their interactions and their location (smart card, untrusted terminal and SLR), depending on the security requirement attached to each of them. The user's profile and the access control engine have to be embedded in the smart card to benefit from its tamper-resistance while the remaining part of the architecture can run in an untrusted environment. The figure shows a particular implementation of MobiDiQ in a smart phone context (*i.e.*, a cell phone equipped with a smart card, typically a SIM card). If the location of the components depends on the targeted platform (for tamper-resistance concerns), the architecture applies the same way to any smart card-enabled device. The figure is self-explanatory considering the technical details given in the preceding sections.

Figure 4 MobiDiQ architecture



From the application developer's point of view, all of the complexity of the internal storage, access control evaluation and security management is confined to the smart card and smart card proxy codes, so that the application developer can concentrate on the application logic.

6 Security and performance issues

6.1 MobiDiQ security

The security issue is twofold. From the content provider and content distributor point of view, security means enforcing the tamper-resistance of licence management. This refers to the licence itself and the data linked to the licence, namely the resource and the contextual information (user's profile and historical data) with which conditions can be defined. From the consumer point of view, security means enforcing the privacy of one's data (profile, historical data). The smart card serves both as a tamper-resistant and private store for contextual information and as an SOE to evaluate and enforce complex business rules based on these contextual data. The information stored outside the smart card (*i.e.*, licences, profile information and historical data stored in the SLR) are kept encrypted and signed to enforce their privacy and integrity. It is worth noting that all communications are encrypted and that a PKI infrastructure is assumed.

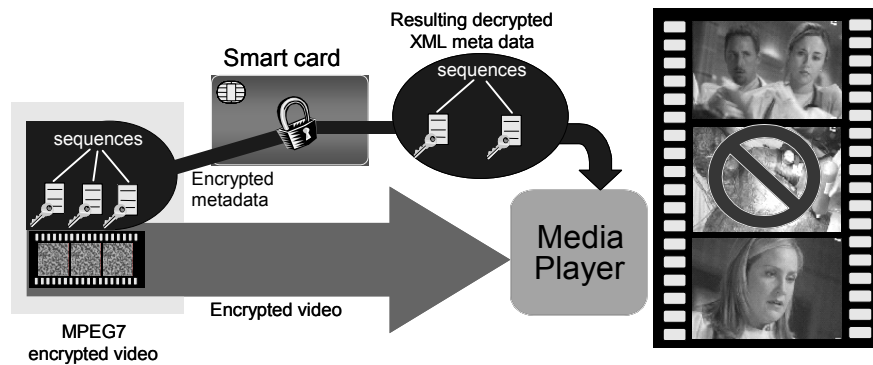
The security level provided by MobiDiQ is bounded by the security level provided by the smart card itself. While smart cards cannot be considered unbreakable devices, they satisfy the highest certification level (EAL 7) of the Common Criteria (2005), the international standard (ISO/IEC 15408) for computer security. One may also question the generality of the MobiDiQ approach, considering that any rendering device is assumed to be equipped with a smart card. Smart cards and other forms of secure chips are today plugged or embedded into a growing variety of devices (*e.g.*, PC, PDA, cellular phone, set-top-box) to serve different applications (certification, authentication, electronic voting, e-payment, healthcare, anti-piracy certification techniques, *etc.*). Thus, SOEs have become a reality on client devices (TCPA, 2006).

6.2 Performance issues and experiments

When dealing with large digital content (*e.g.*, video, music), the traditional performance/security trade-off must be addressed. For maximum security, the whole digital content should traverse the smart card and be decrypted on the fly. This requires a high-bandwidth smart card equipped with a powerful CPU. While some advanced smart card platforms support a full-speed USB throughput, most existing smart cards actually exhibit a much lower communication throughput.

A prototype of MobiDiQ has been developed on such a low throughput SIM card and has been the recipient of the Gold Award of the SIMagine'2005 international software contest (Bouganim *et al.*, 2005). This prototype dealt with a selective dissemination of multimedia streams through unsecured channels. We considered videos encoded using the MPEG7 standard, which allowed the storing of short descriptions of the scenes in the XML metadata. Video objects being quite large and the response time being an important requirement, the smart card used in this experiment was unable to solve the equation owing to its limited communication throughput. To tackle this issue, we traded security for performance as follows. We separated the metadata from the video stream and encrypted the video stream thanks to secret keys stored in that metadata. In this setting, the smart card performs the access control on the XML metadata and delivers the decryption keys to the media player, according to the user's privileges, in the spirit of existing Conditional Access methods (see Figure 5).

Figure 5 MobiDiQ on low-bandwidth smart cards



7 Security and performance issues

Digital piracy is threatening the global multimedia content industry while existing DRM models badly adapt to several new, attractive usage scenarios and exasperate consumers owing to too coercive methods. MobiDiQ aims at reconciling the content providers' and consumers' points of view by giving the ability to develop fair business models (user friendliness, fair superdistribution, ethic enforcement, privacy preservation). In this respect, MobiDiQ can be considered a tamper-resistant enabler of ethical Service Level Agreements.

From a technical point of view, MobiDiQ is an XML-based tamper-resistant right-management engine. The support of XML metadata makes MobiDiQ agnostic about the type of multimedia content to be protected, while the support of XPath access control rules makes MobiDiQ agnostic about the DRM model to be used at the application level. Roughly speaking, the MobiDiQ engine can be seen as a DRM virtual machine with XPath access control rules as bytecode.

A prototype of MobiDiQ has been developed on a SIM card and has been the recipient of the Gold Award of the SIMagine'2005 international software contest. The demonstration of this prototype showed how licences are expressed, exchanged and evaluated in a MobiDiQ-enabled cell phone when playing a video sequence under strict parental control (violent scenes are withdrawn from the video stream).

References

- Bertino, E., Braum, M., Castano, S., Ferrari, E. and Mesiti, M. (2000) 'AuthorX: a Java-based system for XML data protection', *IFIP Working Conference on Databases Security*.
- Bertino, E., Castano, S. and Ferrari, E. (2001) 'Securing XML documents with Author-X', *IEEE Internet Computing*, Vol. 5, No. 3, pp.21–31.
- Bouganim, L., Dang Ngoc, F. and Pucheral, P. (2004) 'Client-based access control management for XML documents', *30th International Conference on Very Large Data Bases (VLDB)*.
- Bouganim, L., Dieu, N. and Pucheral, P. (2005) 'MobiDiQ: mobile digital quietude', *Gold Award of the SIMagine 2005 International Contest*, organised by Sun, Axalto and Samsung.
- Champeau, G. (2004) *Fnacmusic.com: le test complet sur Ratiatum.com – Le Peer-to-Peer (P2P) au delà du téléchargement*, (in French), http://www.ratiatum.com/p2p.php?id_dossier=1708&page=1.
- Chan, C., Felber, P., Garofalakis, M. and Rastogi, R. (2002) 'Efficient filtering of XML documents with Xpath expressions', *International Conference on Data Engineering (ICDE)*.
- Common Criteria (2005) *Common Criteria for Information Technology Security Evaluation Part 1: Introduction and General Model*, August 1999, version 2.3. CCIMB-2005-08-001.
- ContentGuard (2004) *MPEG-21 Right Expression Language (MPEG-REL)*, ISO/IEC 21000-5:2004 standard, <http://www.contentguard.com/>.
- Damiani, E., De Capitani di Vimercati, S., Paraboschi, S. and Samarati, P. (2002) 'A fine-grained access control system for XML documents', *ACM TISSEC*, Vol. 5, No. 2.
- Diao, Y. and Franklin, M. (2003) 'High-performance XML filtering: an overview of Yfilter', *International Conference on Data Engineering (ICDE)*.
- Digital Millennium Copyright Act (DMCA) (1998) *The Digital Millennium Copyright Act (DMCA)*, <http://www.copyright.gov/legislation/dmca.pdf>.
- DRM Watch Staff (2004) *DRM Watch: Fnac Stretches Fair Use Rules in New Online Music Service*, <http://www.drmmwatch.com/oct/article.php/3412101>.

- Electronic Frontier Foundation (2004) *Unintended Consequences: Five Years Under the DMCA*, <http://www.eff.org/IP/DMCA/>.
- European Union Copyright Directive (EUCD) (2001) *European Union Copyright Directive (EUCD)*, <http://www.fipr.org/copyright/eucd.html>.
- Gabillon, A. and Bruno, E. (2001) 'Regulating access to XML documents', *IFIP Conference on Database and Application Security*.
- Green, T., Micklau, G., Onizuka, M. and Suciu D. (2003) 'Processing XML streams with deterministic automata', *International Conference on Database Theory (ICDT)*.
- Hopcroft, J. and Ullman, J. (1979) *Introduction to Automata Theory, Languages and Computation*, Addison-Wesley.
- Inspired (2006) *Inspired, The Future of Smart Card!*, <http://www.inspiredproject.com/>.
- International Federation of Phonographic Industry (IFPI) (2003) *International Federation of Phonographic Industry (IFPI)*, <http://www.ifpi.org/>.
- Miklau, G. and Suciu, D. (2002) 'Containment and equivalence for an XPath fragment', *ACM Symposium on Principles of Database Systems (ACM PODS)*.
- Open Digital Rights Language (ODRL) (2006) *The Open Digital Rights Language Initiative*, <http://odrl.net/>.
- Open Mobile Alliance (OMA) (2006) *Open Mobile Alliance*, <http://www.openmobilealliance.org/>.
- Peng, F. and Chawathe, S. (2003) 'XPath queries on streaming data', *International Conference on Management of Data, ACM SIGMOD*.
- SAX (2004) *Simple API for XML*, <http://www.saxproject.org/>.
- Trusted Computing Platform Alliance (TCPA) (2006) *Trusted Computing Platform Alliance*, <http://www.trustedcomputing.org/>.
- Tual, J.P. (2004) 'DRM as new business driver for smart-card: potential reality or inaccessible dream?', *E-smart*.
- US Department of State (2006) *Fair Use Legal Definition*, <http://usinfo.state.gov/topical/econ/ipr/ipr-glossary.htm>.
- W3C (1999) *The XML Path Language XPath*, <http://www.w3.org/TR/xpath>.
- W3C (2006) *W3C Consortium, 'PICS: Platform for Internet Content Selection'*, <http://www.w3.org/PICS>.
- Westin, A.F. (1999) *IBM-Harris Multinational Consumer Privacy Survey*, <http://www.pco.org.hk/english/infocentre/files/westin.doc>.
- eXtensible Access Control Markup Language (XACML) (2006) *OASIS eXtensible Access Control Markup Language, (XACML)*, <http://www.oasis-open.org/committees/xacml/>.
- eXtensible Media Commerce Language (XMCL) (2001) *XMCL – The eXtensible Media Commerce Language*, <http://www.xmlcl.org/specification.html>.
- eXtensible Rights Markup Language (XrML) (2006) *XrML eXtensible Rights Markup Language*, <http://www.xrml.org/>.

Notes

- 1 Models allowing more than one group could be envisioned to handle specific situations.
- 2 Lucie cannot erase this record from the SLR thanks to cryptographic integrity checking and she cannot remove it from her SIM profile thanks to the tamper-resistance feature of the card.
- 3 Parental control fields are signed using the parental control private key to ensure their authenticity.
- 4 For example, to prevent the child from watching videos during classes or in the school area.
- 5 The metadata have been simplified for space considerations.