Taylor&Francis healthsciences

Security of medical multimedia

S. TZELEPI*, G. PANGALOS and G. NIKOLACOPOULOU
Informatics Laboratory, Computers Division, General Department, Faculty of Technology, Aristotelian University, Thessaloniki 54006, Greece

Abstract. The application of information technology to health care has generated growing concern about the privacy and security of medical information. Furthermore, data and communication security requirements in the field of multimedia are higher. In this paper we describe firstly the most important security requirements that must be fulfilled by multimedia medical data, and the security measures used to satisfy these requirements. These security measures are based mainly on modern cryptographic and watermarking mechanisms as well as on security infrastructures. The objective of our work is to complete this picture, exploiting the capabilities of multimedia medical data to define and implement an authorization model for regulating access to the data. In this paper we describe an extended role-based access control model by considering, within the specification of the role–permission relationship phase, the constraints that must be satisfied in order for the holders of the permission to use those permissions. The use of constraints allows role-based access control to be tailored to specifiy very fine-grained and flexible content-, context- and time-based access control policies. Other restrictions, such as role entry restriction also can be captured. Finally, the description of system architecture for a secure DBMS is presented.

Keywords: Database management systems

1. Introduction

The application of information technology to health care has generated growing concern about the privacy and security of medical information. As health care organizations collect, process, and store more health information in computerized form and use both private and public telecommunications systems to transmit this information between different entities, they must ensure that adequate mechanisms are in place to protect the information. Furthermore, data security and communication security in the field of multimedia are higher [1]. This is mainly due to the complexity of multimedia medical data and their applications. Most of the work on multimedia medical data security until now has focused on cryptographic and watermarking approaches [2–6]. While valuable, cryptography and watermarking technologies are not enough to control access to multimedia medical data. Cryptography and watermarking can only control secrecy, authentication and copyright protection aspects, but cannot handle for example different types of access by different users, content- and context-based access control to multimedia medical data, and fine-grained restrictions

*Author for correspondence; e-mail: tzelepi@eng.auth.gr, gip@eng.auth.gr

Medical Informatics & The Internet in Medicine ISSN 1463-9238 print/ISSN 1464-5238 online © 2002 Taylor and Francis Ltd
http://www.tandf.co.uk/journals
DOI: 10.1080/14639230210153730

at the level of individual users and specific multimedia medical data [7]. Therefore additional approaches should be applied at a higher level.

Various access control policies are needed to ensure that users access only the multimedia medical data for which they have authorization. Although several efforts have been made to develop access control models for traditional database environments, the multimedia medical database environment is quite different to traditional database environments for which such models have been specified. The main difference is that multimedia medical DBMS requires the ability to specify more expressive access control policies than those provided in traditional DBMS. In conventional medical database environments access control is usually performed against a set of role-based authorizations [8], since roles accurately describe which types of people need access to certain types of objects. Such simple role-based authorizations are not well suited for multimedia medical database environments; these require revisiting of such role-based authorizations.

The notion of roles is an important factor in authorization rules, but in order to be effective in a multimedia medical image database system context it has to be used in conjunction with the following information:

- semantic content of the images;
- domain of the health system worked for by a particular caregiver;
- location from where the user is accessing information services;
- patient–physician relationship; and
- time.

Unfortunately, a role-based access control (RBAC) cannot be used to handle the above requirements [9]. In order to overcome this problem, in this paper we propose an extended RBAC model by considering, within the specification of the role–permission relationship phase, the constraints that must be satisfied in order for the holders of the permission to use those permissions. The proposed access control model preserves the advantages of scaleable security administration that RBAC-style models offer and yet offers the flexibility to specify complex fine-grained access restrictions based on the semantic content of the images, the attributes of the user accessing the image, the relationship between the user and the patient whose images are to be accessed and the time. In addition, in our model, by extending the semantics of the constraints in role–permission relationship, other kinds of restrictions such as role entry restriction can also be handled.

A subset of Object Constraint Language (OCL) [10] is used for specifying constraints. In the development of content-based constraints a simplified medical image model for describing the semantic content of a medical image is used. The medical image can be viewed as pairs of iso-semantic regions and signals in respect of an anatomic and a pathological model [11]. Moreover, medical images are associated with complementary textual patient information.

The rest of this paper is organized as follows. § 2 deals with the most important security requirements of today's multimedia medical data. Security measures and security mechanisms which are fulfilling these requirements are discussed in § 3 and 4 respectively. § 5 discusses problems related to the specification of an access control model for multimedia medical data. § 6 introduces related work and contrasts it with our work, while § 7 introduces the medical image data model. § 8 presents an access control model for multimedia medical data. § 9 introduces the access control architecture. Finally, § 10 concludes the paper.

2. General security requirements

A review of the intended use of multimedia medical data indicates that several requirements regarding information security and integrity must be satisfied [5, 6, 12]. These requirements are as follows:

- *Limited access to multimedia medical data* is indicated for several reasons. Patient privacy, medical ethics, and simple good business sense are all important considerations. Patients have an expectation that their medical records and information will be, in general, kept in confidence (academic or other routine use in general are exceptions at least tacitly accepted by most patients). This expectation may, in some jurisdictions, be codified in law; in any case, lawsuits for failure to preserve the confidentiality of medical information are a real threat. Moreover, even if there are no legal ramifications, a hospital or clinic that gains a reputation of having a cavalier attitude towards, or being careless with, the personal information of its patients runs the real risk of losing business to competitors. These factors all argue for restriction of multimedia medical data to only those requiring access.
- *Protection of data against unauthorized disclosure* is also required for several reasons. As is the case with limiting access to multimedia medical data, issues of privacy, medical ethics, and good business sense are all considerations in this case. However, certain categories of multimedia medical data require a greater degree of protection than others. Specifically, patient identification data and the binding of this patient identification data to medical images appear to require a greater degree of protection against unauthorized disclosure than do the images themselves.
- *Positive detection of data corruption* is required for reasons of patient safety. If a medical image or stored patient information has been corrupted, either by chance accident or malicious alteration, it must be detectable. Otherwise, it is possible that a patient will receive inappropriate and potentially fatal medical treatment.
- *Positive binding between patient data and other data* is also required for reasons of patient safety. Medical images and related consultations are often the basis for decisions regarding surgery or other potentially health- or life-threatening medical procedures. There must be a mechanism to ensure a positive, detectable binding between patient identification information (i.e. *of whom* is the image) and the other information (e.g. images, consultation transcripts, annotations) stored on the system. Otherwise, the potential exists for patients to receive medical care that is based on the medical image of someone else, with potentially serious or fatal results.
- *Protection against data corruption* is required for both economic and patient health and safety reasons. While it is preferable to re-image a patient if an image has become corrupted, this should be avoided if possible. Re-imaging a patient is costly, occupies diagnostic equipment, and is inconvenient for the patient.

- *Protection of information against unauthorized distribution and illegal copies* is required to protect intellectual property and content continuity from unauthorized use, misappropriation and misrepresentation.
- The desirability of an *audit trail* is essentially an internal multimedia medical database system issue. A positive audit trail of system resource usage, patient information access, etc, is a key element in providing the other requirements detailed above. It is also invaluable in determining the extent of possible damage and the critical details (how, what, who) in the event that a security incident occurs.

3. Security services

Security services are abstract system or network services that provide certain protections for multimedia medical information being stored or transferred. This section discusses the security services required in multimedia medical database systems security and medical image communications security [5, 6, 12, 13].

3.1. *Multimedia medical database security services*

Access control is essential for multimedia medical database systems. The multimedia medical database system must store (or have direct links to) patient data as well as medical images. Legal and ethical considerations dictate that access to this data must be strictly controlled. *Access control services* generally involve a combination of some form of *peer entity identification* (to authenticate the entity requesting resource usage) with a *rules-based permission system.* An entity requesting access to a system or system resource is generally first authenticated (i.e. its identity is verified). After authentication, a rules-based process determines whether or not this entity is allowed access to the requested resource. It may also require the use of other security services (e.g. confidentiality, data integrity or non-repudiation) when invoked remotely. *Confidentiality services* are intended to protect information against unauthorized disclosure.

Medical image database systems require the *data origin authentication service* in order to ensure that only information from allowable sources is placed into the medical image database system. Otherwise, it could be possible for a malicious outsider to insert apparently valid information into the system.

Mechanisms for *non-repudiation* and *copyright protection* are also required for multimedia medical database systems. This is particularly true for data that relates to patient images. Since medical image databases must serve as long-term repositories for images, they must accept responsibility for these images when received from an authorized source; the medical image database should not be able to later deny receiving this information, nor should the sending entity be able to deny sending it. Additionally, the medical image database should not subsequently be able to forge or distribute images without authorization of this information. The same considerations apply to information stored in databases associated with, but not directly a part of, the image database.

Data integrity and *strong information* binding are each essential for information stored in multimedia medical image databases. Changes to stored information must be positively detectable; the binding between patient information, images, and other related data must be as strong in storage as when in transit. The problem of data integrity in image databases is of particular concern, as these facilities may need to provide long-term protection for stored information and transient protection during transfer.

Availability of information is also a major consideration. The information stored in medical image databases and related databases requires protection against equipment failure, and provisions for continuity of operations in the event of major incidents should be considered.

Administrative and *management policies* must also be considered when determining multimedia medical database security. In general, security policies for multimedia medical database systems should be based on the principles of least privilege and authorized actions. This is necessary in order to limit the potential for damage in the event of both unauthorized access to multimedia medical database systems as well as that posed by a rogue user of the system.

Personnel management issues must also be addressed. The process of screening and hiring users of multimedia medical data must be sufficient to weed out obviously unqualified or untrustworthy candidates. Training in system use and security responsibilities must be provided for persons using multimedia medical database systems.

Finally, issues of *data reliability* must be addressed. Most of the information in a multimedia medical database system is literally irreplaceable; portions of a multimedia medical database system must store data for the long term. Attention to the potential problems associated with long-term storage and periodic backup of on-line data is a must.

3.2. *Medical image communications security services*

Medical image communications security requires the application of several security services in order to be effective.

Multimedia medical data communications security requires the use of *peer entity authentication*. Prior to transferring multimedia medical data, users of the system must be confident that the other entity participating in the transfer of data is in fact who it claims to be. This authentication should be mutual. It is just as important for a remote radiologist to know with certainty that the system from which he or she has requested download of patient case files is in fact the image archive as it is for the image archive to determine that the remote radiologist is in fact an authorized user of the system.

Confidentiality is also required. The assumption is made that at least some portion of the communications pathway involved in the transfer of multimedia medical information to and from remote sites are accessible to persons not having authority to access the multimedia medical database system. This in general is true any time public communications networks are used to transfer multimedia medical information – including those cases in which leased private circuits are used. Physical or electronic access to the communications pathway allows the possibility of unauthorized interception of information, with resulting unauthorized disclosure.

Data integrity services is required for medical image communications security. These services are required to detect the occasional errors in transmission, which occur during transfer of data, as well as malicious attempts to alter data while in transit.

It is also necessary that *information binding* is maintained during multimedia medical database system communications. Proper data integrity, coupled with a sufficient, exportable method of attaining and maintaining strong information binding for multimedia medical data, will ensure that this requirement is met.

With the rapid growth of broadband networks, dissemination and distribution of multimedia medical data among the health care professional community via the Internet is a must way to go. Digital medical images, including still photography and video, which are stored and transmitted within medical communities, are being exposed to a broad audience due to widespread connection of PCs to the Internet. *Multimedia content protection* is necessary to ensure the validity of owner identification, proof of ownership, authentication, covert communication, transactional watermarks (fingerprinting) and copy control of medical images.

4. Providing security services

As noted above, in order to satisfy the general security requirements of multimedia medical data a number of security services are required. These security services include some form of confidentiality, data integrity, access control, non-repudiation, and information binding. This section of this paper will investigate possible methods of providing the requisite security systems in the multimedia medical environment [5, 6, 12, 13].

4.1. *Providing confidentiality*

Effective confidentiality of information while in storage or in transit requires the use of a cryptographic system. Prior to transferring multimedia medical data a wavelet-based technique for detecting and eliminating text in an image is required for confidentiality purposes. Strong access control coupled with effective personnel and security management policies can also be used to control access to information on stand-alone systems.

4.2. *Providing data integrity*

Data integrity in multimedia medical database systems can also be achieved by using various methods. A combination of durable, write-once media, checksumming codes (e.g. CRCs and/or hashing functions, either keyed or unkeyed) and procedural methods (access control coupled with personnel/ security management) can be used to provide effective data integrity for multimedia medical database. Data integrity also requires the use of non-repudiation services. Additionally, embedding extra information into an image using data hiding technology is an effective method for image integrity (e.g. authentication and tamper-proofing).

4.3. *Providing access control*

Provision of access control in multimedia medical database systems involves peer entity authentication. The ITU-T (formerly CCITT) X.509 three-way strong authentication method, coupled with a password protected smart-card type physical token, is recommended for access control. Various other specific access control mechanisms, e.g. an access control base using time of day restrictions, semantic content of images, and medical context, as criteria for allowing/denying/terminating access to date, are also required. If done remotely, confidentiality services (to preclude eavesdropping) are required.

4.4. *Providing non-repudiation*

Notarization, cryptography, or some combination of these features, can provide some level of non-repudiation. The optimal solution is to use a special form of notarization, the X.509 certification hierarchy, coupled with asymmetric keying and unkeyed hashing functions.

4.5. *Providing information binding*

Information binding can be provided either by a combination of confidentiality and data integrity services through the digital signature process or by watermarking technology. By using watermarking technology a permanent association between patient information and medical images is created. For example, by embedding patient information inside medical images, the two become a single entity, thus eliminating both the need to preserve a link between the two types of data and the risk of their separation.

4.6. *Providing multimedia content protection*

Multimedia content protection can be provided by a combination of cryptography and watermarking technology. In particular, watermarking has been considered to be a key technology for multimedia content protection.

Digital watermarking is an important factor in the protection of patient medical images and digital content in the coming years. For example, we can embed encrypted image and patient information into an image. The embedded information can be extracted and decrypted by the receiving site to verify the patient identification and confirm image authenticity and integrity. In addition, the mechanism of embedding patient information within medical images promotes patient safety and record consistency. One area that will enable digital watermarking to be considered a successful application in the medical field is when the alteration of a watermarked image becomes detectable to the point that attempting to remove a watermark from its object would render that object useless.

5. Access control policy requirements

Access control policy is needed to ensure that users access only the multimedia medical data for which they have authorization. Although several efforts have been made to develop access control models for traditional database environments, the multimedia medical database environment is quite different to traditional database environments for which such models have been specified. The main difference is that multimedia medical DBMS requires the ability to specify access more expressive control policies than those provided in traditional DBMS. In conventional medical database environments access control is usually performed against a set of authorizations stated by security officers. An authorization in general is specified on the basis of three parameters $<u, o, p>$. This triplet specifies that user u is authorized to exercise privilege p on object o. Such simple paradigm is not well suited for multimedia medical database environments; these require revisiting of such paradigms. Furthermore, in a multimedia medical database environment the resources to be protected are not only traditional data but also multimedia. Such peculiarities call for more flexibility in specifying access control policies. The access control mechanism must be flexible enough to support a wide spectrum of *heterogeneous protection objects*.

A second related requirement is the support for *content-based access control*. Content-based access control allows one to express access control policies that take the protection object content into account. This is an important requirement since very often protection objects with the same type and structure have contents of different sensitivity degrees. For example, all images presenting a cancer of the lung must not be made available to physicians who are accessing information from no trust domain. In order to support content-based access control, access control policies must allow one to include conditions against protection object content.

A third requirement is the support for *context-based access control*. Context-based access control allows one to express access control policies based on the domain of the health system a particular caregiver works for, on the location where the user is accessing information services from, on the patient–physician relationship and on the time.

A fourth requirement is related to the *heterogeneity of subjects* which requires access control policies based on user responsibilities and qualifications, rather than based on very specific and individual characteristics (e.g. user IDs). A possible solution, which better takes into account user responsibilities and qualifications in the formulation of access control policies, is to support the notion of role. A role can be seen as a set of actions or responsibilities associated with a particular working activity. Under role-based access control models, all authorizations needed to perform a certain activity are granted to the role associated with that activity, rather than being granted directly to users. Users are then made members of roles, thereby acquiring the role's authorizations.

Finally there is a need to support *fine-grained authorizations* at the level of individual users and specific images. For example, just because the doctor's role enables a set of accesses to medical images does not mean that the doctor's role should provide access to all medical images. A doctor can only access the medical images for those patients currently assigned to this doctor.

6. Related work

One of the problems of applying RBAC to multimedia medical image database systems is the specification and enforcement of fine-grained access control at the level of individual users and specific images. As mentioned above, just because the doctor's role enables a set of accesses to medical images does not mean that the doctor's role should provide access to all medical images. A doctor can only access the medical images for those patients currently assigned to this doctor. Unfortunately RBAC cannot meet completely this requirement [9]. There have been several approaches for creating an instance level policy for roles by using the notion of team-based access control (TMAC) [9] or by introducing parameterized roles to RBAC models [14, 15].

An alternative approach to specifying and enforcing fine-grained access control at the level of individual users and specific images is proposed in this paper by considering, within the specification of the role–permission relationship phase, the constraints that must be satisfied in order for the holders of the permission to use those permissions. Furthermore, constraints offer the ability to specific complex access restrictions based on the semantic content of the images, the attributes of the user accessing the image, the relationship between the user and the patient whose images are to be accessed, and the time.

Constraints have also been addressed in [15], where content-based access control was enforced simply by specifying some constraints against attribute values of data objects. In contrast, due to the nature of medical images, content-dependent access control for a medical image database system must be based on the semantics of the medical images, rather than on the attributes characterizing them. Medical image attributes often only deal with physical characteristics of the medical images (for example, acquisition device, direction, format) and therefore are not significant for access control.

In order to develop content-based constraints specification, a simplified medical image model for describing the semantic content of a medical image is described in the next section.

7. The underlying medical image data model

For content-based access control, medical image databases must have the capability to recognize and quantitate image content and merge the quantitated image data with textual patient data into a common data model. Established algorithms can readily be integrated into a multimedia medical DBMS for image segmentation, texture analysis, content extraction and image registration. These can be performed automatically or interactively depending on the difficulty of segmenting and extracting semantic structures.

The content of a medical image must include three levels of abstraction [17]: (i) features or attributes, such as the texture pattern of an organ, the shape of a breast tumour, and the signal strength of brain metabolites; (ii) structural objects, such as image segments corresponding to anatomical regions; and (iii) semantic information, such as the types of relationships between two objects.

The image features must be further divided into primitive and logical ones. Primitive features are directly obtained from the medical images (e.g. volume, shape, and texture of certain organs in CT images). Logical features are abstract representations of images at various levels of detail and deeper domain semantics (e.g. whether the volume of an anatomic structure is normal). These logical features are synthesized from primitive ones and additional domain knowledge. All features extracted must be entered into appropriate attributes of the objects defined in a data model to facilitate subsequent information access by content.

In this section, we use in the development of our content-based access control model a simplified medical image model which was introduced in [11]. Our emphasis in this paper will be on the development of our content-based access control model rather than on effectively modelling multimedia medical data. In [11], the medical image model represents (see figure 1):

- what is the image (i.e. the pictorial attributes of image);
- what is around the image (i.e. the context); and
- what is in the image (i.e. the content)

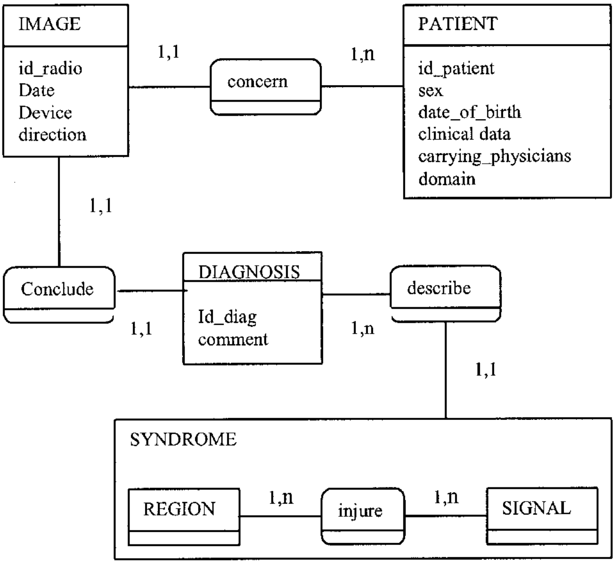

Figure 1. Medical image model.

In [11] a medical image is considered to be an organic structure, where pathological signals can be detected in special places. Therefore, the semantic content of a medical image can be described in terms of iso-semantic regions and signals. In our proposed access control model, the content-based constraints we are concerned with deal with the presence or absence (and eventually with the characteristics) of a referenced object in a particular location of the image. For example, based on the above medical image data model the following content-based constraint denotes all images x that present a tumour with a height of more than 1 cm on the left ventricle.

$$\text{'left ventricle'} \in \text{regions(x)} \wedge \text{'tumor'} \in \text{signals}(x, \text{'left ventricle'}) \wedge \text{height}(x, \text{'left ventricle'}, \text{'tumor'}) \geq 1 \text{ cm}$$

8. An extended role-based access control model for multimedia medical databases

The basic components of a simplified RBAC model are users, roles and permissions [17]. The user is a person who uses the system or an application program in the system. Membership to roles is granted to users based on their obligation and responsibility in the organization. The operation of a user can be carried out based on the user's role. The role is a set of functional responsibilities within the organization. The system defines roles and assigns them to users. A user–role (U-R) relationship represents collection of a user and a role.

Permission is the way for the role to access to more than one object in the system. The terms authorization, access right and privilege are also used in the literature to denote permission. Permissions are always positive and confer the ability to the holder of the permission to perform some actions in the system.

The simplified RBAC model consists of user–role (U–R) relationship and role–permission (R–P) relationship. A role-permission (R–P) relationship describes which role is assigned to perform what kind of permission in the organization. In this paper, an extended simplified role-based access control model for multimedia medical image database systems is presented. Three major extensions to the model are introduced.

The first extension introduces the notion of *domain*. A domain is a collection of subjects/objects, which have been explicitly grouped together for the purposes of management. A domain will be an organization or part of an organization with a clearly defined professional or statutory role. For example, individual hospitals, private firms offering administrative or quality control services to hospitals and analysis laboratories, are all candidate domains in an interconnected system. The concept of a domain is very similar to that of directory in a typical hierarchical file system. The authorization, which implies to a domain, will, by default, propagate to subdomains and to the subjects/objects within them.

The second extension introduces the notion of *location*. A location is a place from where the user is accessing information services. Location information is used in several types of authorization rules. One type uses location to identify the trust domain from where the user is accessing information services. A reasonable policy would deny access to any sensitive information to anyone attempting to access it from such areas. Location can also be used to derive the emergency level of access. A policy can allow read access to all images of all patients for any user assigned to the role physician and accessing the information from an emergency room.

The third extension concerns the role–permission relationship. In the proposed model, we consider, within the specification of the role–permission relationship phase, the *constraints* that must be satisfied in order for the holders of the permission to use those permissions. In this case, each role–permission relationship is a decision rule, which specifies, besides the access modes the holder *s* of the permission is authorized for on image(s) *i*, also the constraints to be satisfied in order for *s* to exercise the access modes.

In a multimedia medical database context, the general form of a role–permission relationship is quintuple:

$$< identifier,\ s{:}\ r,\ \{action\},\ t{:}\ target,\ constraints(s,\ t) >$$

According to the above definition, a role–permission relationship has the following components:

- identifier: used to identify uniquely the permission
- *s*: the subject to which the permissions apply
- *r*: the role which can process this permission. Subject *s* is authorized for role '*r*'
- action: the operation to be processed by role
- *t*: the object on which actions are to be performed
- target: the object type
- constraint(*s*, *t*): limit the applicability of the permission. Constraints must be satisfied by *s* and *t*.

A subset of Object Constraint Language (OCL) is used for specifying constraints. The use of constraints allows role-based access control to be tailored to specify very fine-grained and flexible content-, context- and time-based authorization policies. For example, consider a health-care organization security policy, where the organization is composed of several hospitals, and each hospital is structured into some divisions. A primary physician is assigned to a division and he/she can only access medical images for those patients currently assigned in that division and to this doctor. In order to achieve such fine-grained policy, we define the following role–permission relationship:

$$\{\text{dp1}, s\text{: Primary\_Physician}, \{\text{view}\}, t\text{: Image, domain\_user(s)} = \text{domain(t)} \wedge s \in \text{carrying\_physicians(t)}\}$$

The function *domain_user* gives us the domain associated with a user. For instance the domain associated to the primary physician 'John Smith' is 'hospital1/div2'. Based on the above policy, 'John Smith' can only access medical images for those patients currently assigned to the domain 'hospital1/div2' and to the doctor 'John Smith'.

In addition, in our model by extending the semantics of the constraints in role–permission relationship other kinds of access restrictions such as role entry (authentication) restriction can also be handled. For example, the following expression states that a user in the roles of Haematologist and Senior Doctor can take on the role of Senior Haematologist.

$$\{\text{dp2}, s\text{: seniorHaematologist } x, \{\}, \{\}, s\text{:Haematologist} \wedge s\text{:SeniorDoctor}\}$$

The values of the above variables will be instantiated at the time of a user's role entry or data access request from information about the user established at logon, or from auxiliary information that is retrieved by database look-up.

9. Implementation

We have implemented a first prototype system for the access control model presented in this paper. The complete system architecture is depicted in figure 2. The components of this system are access control manager, authorization manager, image data manager and image-postprocessing manager. The *access control manager* implements the access control algorithm. Through the *authorization manager*, the security administrator can add, modify, or delete user–role relationships, role–permission relationships and user attributes. The *image data manager* is responsible for handling of images. Each time a new image is acquired by the medical image database system, it is first processed by the *image-postprocessing manager*, which extracts the semantic content from this image. All features extracted are entered into appropriate attributes of the objects defined in the data model to facilitate subsequent information access control by content. Our emphasis at this stage is on those features that have direct impact on the building of a working prototype, rather than on query performance.

This system is based on a two-tier client-server architecture in which the user interface runs on the client and the database is stored on the server. The prototype

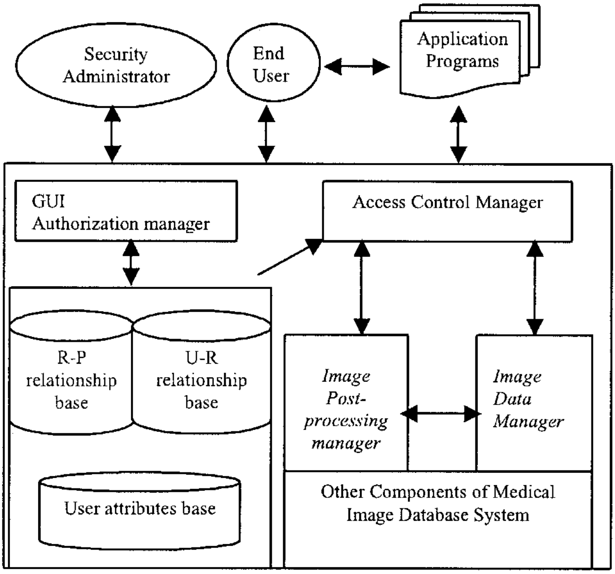

Figure 2. System architecture for a secure multimedia medical database management system.

system has been developed on top of the Oracle 8 DBMS, using Visual C++. We use a relational database management system to implement our image model and to store all necessary data. The connection to the database is accomplished through ODBC. The client initiates the queries and sends the access queries to the database server. After receiving the queries, the server searches and retrieves the relevant results and returns the query results to the client, according to the specified security policy. Then the client displays all the query results. In particular, after receiving the access request, the server verifies whether user u, trying to access image i at the time t from location l, belonging to domain d, using a privilege p, under a certain role r, is authorized to do so, according to access control restrictions enforced by that role.

The client of the system plays two important roles. One role is to interact with users and the other role is to communicate with the server. Since the client is the front-end of the system, graphical user interfaces that are easy-to-use, user-friendly, and fully functional have been developed. In addition, the client is structured in the way that future updates can be easily done. The client is broken down into two modules: the query module and the multimedia presentation module.

The task of a multimedia database server is to build up a system that can manage multimedia medical data in an efficient way. The main functions are to access, process, manage and manipulate multimedia medical data. It also serves to filter all client requests through various centrally controlled levels of security, such as access control. Since our authorization model relies on the use of semantic content for supporting content-based access control, our prototype system has been integrated with a mechanism that extracts image features from images. The mechanism we use for image features extraction and semantic objects quantification is based on the work reported in [18] about the detection of brain lesions on SPECT images. Each time a new image is acquired by the medical image database system, it is first processed by the *image-postprocessing manager*, which extracts the semantic content from this image. Information on the semantic content is then stored into Oracle tables and used by our prototype system to perform content-based access control restrictions.

Through the *authorization manager*, the security administrator can add, modify, or delete user–role relationships, role–permission relationships and user attributes. For each administrative operation supported by our model, our prototype system provides a graphical interface that supports the security administrator in performing the administrative operation.

In a multimedia medical database context, low level operations, such as physical read and write operations, are not semantically meaningful for access control in multimedia medical image databases. Therefore, in our model we introduce a set of abstract operations that are relevant to the way users actually access medical images.

Users of medical images database go through the following stages. The user first submits a request for a given image. The medical image database server processes the request, and returns to the user either the annotations associated with the image, or the requested image in thumbnail format. The user can then request the display of the full-resolution image in the main display window. Another group of operations for which access control should also be provided include operations for processing, annotating, deleting the images or for introducing new images. In general, authorizations to perform such operations should be given to few, selected users. For example, saving on the main archive is usually reserved for radiology. Deleting is reserved for the security administrator but only with dual access code security, i.e. someone else with security must be present to authorize the deletion. Retrieval, annotation and processing can be performed by any clinician with the proper access. Technicians and managers need access to management functions. The different modes of operation and image access privileges that are provided as part of our model are described in table 1.

10. Summary

The need for medical image security techniques is growing with the increasing use of digital techniques for transmitting and storing image. Generally, there exist two kinds of technique for image data protection: cryptography and digital watermarking. In order to complete this picture, here we define and implement an extended role-based access control model by considering, within the specification of the parametric role–permission relationship phase, the

Table 1. Image privileges provided by the access control model.

Privilege	Meaning
View_annotation	To display the results as the associated annotations only.
View_thumbnail	To display the requested image in thumbnail format. This speeds up the query response time.
Display	To display the full-resolution image in main display window.
Edit_annotation	To edit the annotations of the images.
Edit_image	To process, delete or add images to the medical images database.

constraints that must be satisfied in order for the holders of the permission to use those permissions.

The novelty of our model is its support for a set of features, such as content-based authorizations on multimedia medical data and the specification of authorizations based on the role of the user, the domain of the health system for which a particular caregiver works, the location from where the user is accessing information services, the patient–physician relationship, and the time, which are crucial for a health care environment. In this paper we have addressed how to build image access control based on both image semantic contents and security context, which is a more flexible and natural way to express authorization for image data. By security context, we mean the role of the user, the domain of the health system for which a particular caregiver works, the location from where the user is accessing information services, the patient–physician relationship, and the time.

The use of constraints allows role-based access control to be tailored to specify very fine-grained and flexible content-, context- and time-based access control policies. In our model, by extending the semantics of the constraints in role–permission relationship, other types of restrictions such as role entry restrictions can also be handled. The use of parametric role–permission relationships allows our model to denote a set of authorizations in a compact way. In addition, the proposed access control model preserves the advantages of scaleable security administration that RBAC-style models offer. Furthermore, our access control model can be coupled with any multimedia system able to associate information in forms of semantic contents to different multimedia data.

References

1. Wohlmacher, P., 1998, Requirements and mechanisms of IT-security including aspects of multimedia security. *Proceedings of the Multimedia and Security Workshop at ACM Multimedia '98*, Bristol, UK, September 1998, edited by S. Dittmann, P. Wohlmocher, P. Horster, R. Steinmetz, pp. 11–19.
2. Smith, J. P., 1995, Authentication of digital medical images with digital signature technology. *Radiology*, **194**, 771–774.
3. Wong, S. T. C., 1996, A cryptologic-based trust center for authenticating medical images. *J. American Medical Informatics Assoc.*, **3**, 410–421.
4. Wolfgang, R. B. and Delp, E. J., 1997, Overview of image security techniques with applications in multimedia systems. *Proceedings of the SPIE Conference on Multimedia Networks: Security, Displays, Terminals and Gateways*, **3228**, 2–5 November 1997, Dallas, Texas, pp. 297–308.
5. Macq, B. and Dewey, F., 1999, Trusted headers for medical images. *Proceedings of the DFG $V^{III}D^{II}$ Watermarking Workshop 1999*, 5–6 October 1999, Erlangen, Germany.
6. Berry Rit, T., 2000, Digital watermarking and medical imaging. Cryptography Course homepage, 8 November 2000. Available at http://citeseer.nj.com/406247.html.
7. Fernandez, E. B. and Nair, K. R., 1998, An abstract authorization system for the Internet. *Proceedings of the 9th International Workshop on Database and Expert Systems Applications* (DEXA '98) (Vienna: IEEE Computer Society) pp. 310–315.
8. Mavridis, I., Pangalos, G. and Khair, M., 1999, eMEDAC: Role-based access control supporting discretionary and mandatory features. *Proceedings of the 13th IFIP WG 11.3 Working Conference on Database Security*, Seattle, Washington, USA, edited by V. Atluri and J. Hale (Kluwer Academic Publishers), pp. 63–78.
9. Thomas, R. K., 1997, Team-based access control (TMAC): A primitive for applying role-based access controls in collaborative environments. *Proceedings of the ACM RBAC '97*, Fairfax, VA, USA, pp. 13–19.
10. Rational Software Corporation, 1997, Object Constraint Language Specification, Version 1.1, Available at http://www.rational.com/uml/, September 1997.

11. Tchounikine, A., 1997, Creation and content-based retrieval in a radiological documentary record. *Proceedings of the 3rd Basque International Workshop on Information Technology*, Biarritz, France, July 2–4 (IEEE Press) pp. 111–117.
12. Medenis, M. M., 1997, Security in teleradiology systems: requirements and proposed mechanisms. Submitted to Dr Ralph Martinez, Electrical and Computer Engineering Department, University of Arizona for *ECE678*, Integrated Telecommunications Networks Spring Semester, 1997.
13. Wolfgang, R. B. and Delp, E. J., 1997, Overview of image security techniques with applications in multimedia systems. *Proceedings of the SPIE Conference on Multimedia Networks: Security, Displays, Terminals and Gateways*, **3228**, 2–5 November 1997, Dallas, Texas, pp. 297–3308.
14. Giuri, L. and Iglio, P., 1997, Role templates for content-based access control. *Proceedings of the Second ACM Role-Based Access Control Workshop*, November 1997, Fairfax, VA, USA, pp. 153–159.
15. Lupu, E. C. and Sloman, M., 1997, Reconciling role-based management and role-based access control. *Proceedings of the Second ACM Role-Based Access Control Workshop*, November 1997, Fairfax, VA, USA, pp. 135–141.
16. Wong, S. T. C. and Huang, H. K., 1996, Design methods and architectural issues of integrated medical image data based systems. *Computerized Medical Imaging and Graphics*, **20**, 285–299.
17. Sandhu, R., Coynee, E. J., Feinsteinn, H. L. and Youman, C. E., 1996, Role-based access control models. *IEEE Computer*, **29**, 38–47.
18. Stamatakis, E. A., Glabus, M. F., Wyper, D. J., Barnes, A. and Wilson, J. T. L., 1999, Validation of statistical parametric mapping (SPM) in assessing cerebral lesions: a simulation study. *Neuroimage*, **10**, 397–407.