

Special Issue on "Digital Libraries in Medicine"

An Integrated Architecture for the Provision of Health Telematic Services based on Digital Library Technologies

M. Tsiknakis¹, C. E. Chronaki¹, S. Kapidakis¹, C. Nikolaou^{1,2}, and S. C. Orphanoudakis^{1,2}

¹Institute of Computer Science, FORTH, Heraklion, Crete, Greece

²Computer Science Department, University of Crete, Heraklion, Crete, Greece

Received: May 13, 1997

Revised: June 6, 1997

Abstract. Advances in information technology during the past decade have resulted in a proliferation of clinical information systems dedicated to different user groups and clinical functional areas. This, in turn, has created the need for hospital-wide management and integration of information, and has triggered major efforts towards the development of integrated hospital information systems as a building block of integrated regional healthcare networks.

The diversity of hospital organizations, the complexity of clinical protocols and procedures, as well as the different preferences of various user groups make it extremely difficult for a single *monolithic* information system to effectively serve the needs of an entire healthcare organizational structure. Thus, information and telecommunications systems must primarily provide the infrastructure to support the effective integration of distributed and heterogeneous components, ensuring overall integrity in terms of functional and information interworking. This approach, i.e. the integration of heterogeneous autonomous distributed systems, to developing and managing regional healthcare networks ensures the transfer and integration of consistent information between healthcare facilities, without imposing constraints on the operation of individual clinical units.

This paper presents the results of an ongoing effort for the design and implementation of an architecture, based on digital library technolo-

gies, for the provision of user-oriented telematic services in a regional healthcare network. Specifically, it addresses issues related to the provision of user-oriented services, transparent to the needs of different user groups and the requirements of specific tasks, based on: a) *meta-information* for the creation of an information infrastructure for the regional healthcare network which is, effectively, a multimedia distributed digital library, b) *intelligent information retrieval strategies* to selectively retrieve information from multimedia data, c) *agent-based technologies* for effective service delivery adapted to the current user needs and the task at hand, and d) *middleware services* that explicitly reveal not only the characteristics of the information sources, but also address the context of specific telematic services, through appropriate mediation mechanisms.

Key words: virtual patient record, regional healthcare networks, telematic services, digital medical libraries, metadata, mediation, middleware services, healthcare information infrastructure, agent-based technology.

1. Introduction

The explosive growth in network connectivity and the rapid advances in computing power are replacing the older notion of standalone information utilities with newer notions of interconnected digital libraries. In the digital library, users must simultaneously access multiple distributed information sources that differ in con-

tent, form, and source types. In this context, conventional notions of information retrieval limit the kind of interaction that is observed in actual work. The wide variety of ways in which information is arranged, sought, and used in physical workspaces strongly suggests the importance of providing an equally rich and flexible environment for interacting with information in electronic workspaces [58].

The current trend in healthcare information technology is increasingly digital and multimedia-oriented. The next generation of healthcare information systems will consist of a large number of heterogeneous, autonomous, and distributed information systems, knowledge-intensive applications, and large quantities of multimedia medical data. Hence, as stated in [74] 'a key challenge facing system researchers and builders is to provide a new organizational framework that can integrate this heterogeneous collection of resources into what appears to be a uniform conglomeration of data and knowledge to increase the availability of previously inaccessible information and to address the demanding information processing requirements of modern medical applications'. Many other researchers and relevant bodies are recognizing the requirement for developing a Healthcare Information Infrastructure (HII) [10,12,15,29,44,46,75,76], introducing concepts such as the Medical Image Informatics Infrastructure within the integrated hospital environment [32,74].

The creation of the Healthcare Information Infrastructure (HII) is driven by, among other things, the need for:

- data on outcomes of medical cases that will enable effective choices and compensation of providers,
- automation of mundane tasks to place the focus on patient needs rather than paperwork,
- empowerment of consumers to become more actively involved in their own healthcare,
- flexible remote access to relevant information in order to ensure the continuity of care from the site of an accident, to the healthcare centers and the home of the patient,
- continuous process improvement through integrated and distributed information technology.

Existing and emerging digital library technologies provide a promising approach to the complex problems involved in any effort at developing the healthcare information infrastructure.

In fact, the healthcare information infrastructure may be viewed as a distributed digital library.

This paper presents initial results of an ongoing effort for the design and implementation of an architecture, based on digital library technologies, for the provision of user-oriented telematic services in a regional healthcare network. Specifically, it addresses issues related to the provision of integrated telematic services transparent to the needs of different user groups and the requirements of specific tasks, employing agent-based technology for intelligent information management and retrieval, as well as meta-information and other enabling middleware services. Meta-information is used in the creation of a healthcare information infrastructure model for the regional network (i.e. HRDM), which includes specific components of the virtual patient record, the regional healthcare resources and actors (section 4.2.2). Intelligent information retrieval strategies are employed to selectively retrieve regional information pertinent to the current user and the on-going medical act [2,5,56,71]. Information, interface, and task agents contribute towards effective service delivery, responding to the current needs of each user and the requirements of each task of interest [63] (section 4.2.3). Finally, a range of middleware services ensures the seamless and secure integration of the various information sources [72].

2. Current Trends in Healthcare Telematics

The healthcare organizational structure is naturally distributed, consisting of geographically distributed medical centers in a hierarchy of regional hospitals down to individual general practitioners. The objective of this structure is to offer comprehensive medical care at a local and regional level with continuity across different levels of the hierarchy. Moreover, due to the greater mobility of patients and the population as a whole, national, and international healthcare networks are increasingly used to facilitate the sharing of healthcare-related information among the various healthcare actors. This sharing of information resources is generally accepted as the key to substantial improvements in productivity and better quality of service [65]. Hence, although each healthcare facility is autonomous and devoted to the delivery of a particular set of services, continuity of care requires that different healthcare facilities, offering complementary services or different levels of expertise, ex-

change relevant patient data and operate in a cooperative working environment [16,22].

During a single healthcare episode, medical care is administered by many professionals, involved in a variety of medical acts. Healthcare administration personnel, healthcare professionals, social care professionals, as well as patients need to selectively interact with healthcare-related information. Each of these user groups has different needs in terms of information access, security and quality of service, and is involved in different tasks, medical acts, and healthcare procedures. In addition, the requirements of each user affect information retrieval and presentation strategies, and the overall interaction with the information space.

The evolving healthcare information infrastructure is also characterized by the fact that the bulk of medical data is managed by different specialized systems, designed to support the search for particular data types. For example, MRI scans, EKGs, laboratory reports, and administrative records of the same patient typically reside in separate autonomous information systems. Thus, comprehensive information about a patient is difficult to obtain efficiently, unless these distributed segments of an electronic patient record can be integrated into a virtual patient record and viewed on-line through a unified user interface and visualization environment [20,21,34,40,55,56, 64].

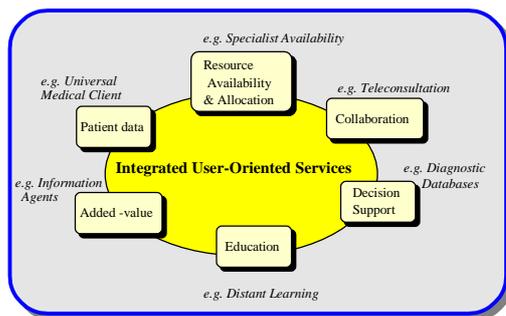


Fig. 1: A regional healthcare network should provide integrated user-oriented services, in the context of a collaborative environment that ensures continuity of care under strict security and authorization policies.

Thus, current trends in healthcare call for *integrated user-oriented telematic services* or end-user services, which ensure prompt and secure access to information resources, provided proper authorization is available. To achieve this goal, complex problems and issues related to data heterogeneity, heterogeneity of platform and service requirements, complex protection

and authorization policies, and interoperability protocols and standards for information exchange need to be addressed. Thus, the healthcare information infrastructure must primarily provide the framework for the effective integration of distributed and heterogeneous components, ensuring overall integrity in terms of functional and information interworking, while advances in network technology should enhance and extend applications, rather than replacing or making them obsolete [45].

2.1 Integrated User-Oriented Telematic Services

A regional healthcare network providing integrated user-oriented telematic services to user groups with different information and quality of service needs presents a significant degree of complexity. This is partly due to the wide spectrum of the services provided. Furthermore, interoperability among telematic services and stand-alone applications is critical, since services and applications share data and information. For example, during a remote teleconsultation session, a specialist in a central hospital who is consulted on the state of a patient, needs to have access to patient data such as patient history, laboratory exams, radiology exams, etc. Moreover, since these telematic services contribute multiple access points to the information space, consistency issues are of major importance. For example, a radiologist within the hospital may use a reporting workstation to enter a diagnostic report, while the referring physician may access the same information through a universal medical browser [20,40,64]. Hence, user-oriented telematic services need to be part of a collaboration environment that ensures the continuity of care and information sharing, under strict security and authorization policies. As shown in Fig. 1, the telematic services which are typically delivered within a regional healthcare network may be classified into five basic classes: *patient data*, *resource availability and allocation*, *collaboration*, *decision support*, *education*, and *added-value services*.

Patient data-oriented services provide access to the critical core of data managed by the regional healthcare network by facilitating alternative views of the virtual patient record, which consists of all signed or certified patient-related information including demographics, reports, billing, etc. As a consequence, they display a high degree of complexity and heterogeneity both in terms of the contributing information sources and the user groups that

need access to these services. Different views of the relevant information space need to be established for different user groups based not only on their information needs and privileges, but also on the current medical act or procedure. Alternative views may be encounter-centered, problem-centered, task-centered, patient-centered, or user-defined temporal snapshots. Furthermore, due to the sensitivity of the information maintained by the component information systems, advanced security and authentication methods need to be employed. In addition, patient data services need to interoperate with other classes of services to enable cooperation among users and access to healthcare resources and facilities.

Resource-oriented services facilitate the interaction of various users with information and services related to resource availability and allocation. User services in this class include facility information, appointment booking, specialist availability, and ambulance assignment. Again, different user groups have access to different views of the healthcare resource information at a different quality of service or priority. For example, the public should have access to information on the facilities and expertise available in each healthcare center and be able to book an appointment with a specific physician or department [8]. As another example, the operator of an emergency coordination center needs instant access to accurate ambulance availability, location, and equipment information to effectively manage a response to an emergency [41]. These examples involve different user groups with different quality of service needs and different authorities.

Collaboration services bridge the gap created by the physical distance separating the users. Note that collaboration is an essential aspect of any service, since 'it promotes social interaction and person-to-person exchange of vital information' [68]. This class of services enables teleconsultation sessions among healthcare professionals to compensate for the lack of experienced or specialized personnel, to address emergency situations, to evaluate the severity of a situation, etc. It effectively makes expertise a shared resource, wherever it may exist. As another example, telemonitoring services establish collaboration between the patient and physician and enable the delivery of healthcare at the home of the patient. Thus, patients receive timely attention, while the operational costs of healthcare centers are reduced. Furthermore, as far as community care is concerned, collaboration among user groups

that share the same chronic medical condition may provide comfort through the sharing of useful information and experience (special interest groups).

Decision support services provide access to the distilled knowledge associated with digital libraries, as well as specific medical procedures and acts. Example services in this class include a drug interaction service provided by a registered pharmaceutical company, access to diagnostic data bases, access to epidemiological studies and reports, as well as evaluation tools and methods which can be provided by special instruments. Usually, these services are provided by specialized systems located at referral healthcare facilities. Despite their importance, these services cannot be afforded by small healthcare centers, even though they may be critical in the treatment of geographically isolated patients. Remote access to these services is usually restricted and charged. These services may be offered through an advanced network of digital medical libraries.

Educational services are especially important in a rural region where access to information sources is scarce. Healthcare professionals should be able to access the educational material of digital medical libraries. At the same time, the public should have access to information relevant to diseases common in their region, so that public awareness is raised and continuity of care is effectively supported. Thus, geographic isolation may not hinder the sharing of knowledge among healthcare professionals, since medical students and healthcare personnel can benefit from the interaction with repositories of medical cases and remote teaching facilities [37,43].

Finally, *added-value services* extend the services discussed in the previous paragraphs and provide specialized support to healthcare professionals. Image processing services, access to specialized indexing and search engines, information filtering and prefetching, content-based image indexing and retrieval engines, and personalized interface, task, and information agents are all examples of services that facilitate efficient and effective utilization of the information space [49,63,66,70].

3. Architectural Framework

Two models of the HII are particularly useful for identifying fundamental technology issues. The HII functional services model [75,76] described below provides a logical model of the relationships between applications and the un-

derlying services that support their development and use. The HII reference architecture model, identifies the major functional elements of the infrastructure and their interfaces and protocols. The latter is the one used in our work and is presented in section 4.1.

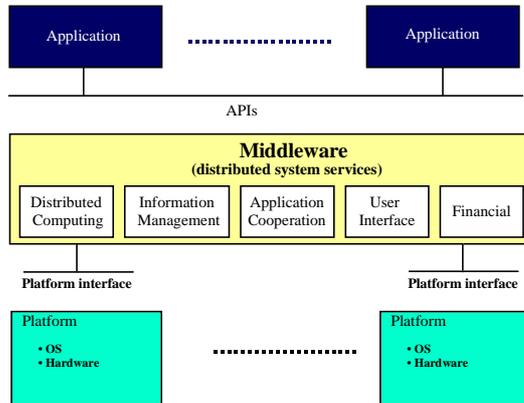


Fig. 2: Layered architecture of the healthcare information infrastructure.

3.1 Healthcare Information Infrastructure: Functional Services Model

Users are primarily interested in information processing applications, which they may own or gain access to as end-users via communications networks. These services are 'enabled' by other underlying, transparent services provided by information and network service providers. Applications and enabling services will in turn employ certain information processing services and systems for data transport, which may be distributed throughout the HII. Thus, the HII consists of three basic components: applications, enabling (or middleware) services, and physical infrastructure. These components can be mapped onto the functional services framework model shown in Fig. 2, as a layered architecture [7].

The functional services model is similar to models of network protocols, but is more generally intended as a way to think about the components of a feature-rich, flexible, open, and distributed infrastructure. The physical infrastructure contains the basic processing and communication components of an information system. The enabling services provide general system-related functionality for applications using the infrastructure. These services are those essential for the HII to meet its goals and requirements and are categorized as follows:

- *Distributed computing services* provide the functionality that links multiple sepa-

rate nodes into one distributed system (network services, location services, security services, etc.)

- *Information management services* organize, store, and retrieve information.
- *Application cooperation services* enable applications to "cooperate" in order to create common multiple end-user services. Typical services include messaging services, object request brokers, workflow managers, agents, and encapsulation facilities.
- *User interface services* present information to and acquire information from the users, thus providing the necessary link between users and the HII.
- *Financial support services* support all commercial and personal financial transactions.

Finally, applications are information processing tools that provide functionality to the user. The underlying enabling services and physical infrastructure supply the means by which applications deliver their functionality.

3.2 Healthcare Information Infrastructure Requirements

The development of a HII for the delivery of customized, secure, integrated user-oriented telematic services requires the design of an architectural framework that is heterogeneous, open, distributed and extensible, with multiple points of access [15]. In developing an integrated framework for the provision of health-telematic services, one must consider the heterogeneity of information sources, user needs and privileges, as well as tasks to be performed. Specifically, such a framework must dynamically encompass new types of applications and systems including real-time monitoring and alarms (e.g. ICU), transaction systems (e.g. billing), office automation tools, and consumer-based interactive systems (e.g. WWW). The employment of *standardized* healthcare components enable plug-n-play integration of distributed healthcare information systems [24].

Specific requirements for the development of the HII are security and authentication, quality of service, integration, extensibility in terms of information sources and services as well as support for task-oriented workflows and multiple overlapping views of the information space.

Different user groups and individual users should acquire different privileges and authorities, in accordance to their current role and the

on-going medical act. The explicit enforcement of appropriate *security* and *authentication* policies is a critical factor in the adoption of new technology in the healthcare domain.

Quality of service is another major requirement, for information to be delivered in a timely manner, without compromising the accuracy of this information. Furthermore, one has to take into account timing requirements in determining the precision and quality of data to be transmitted, so that bandwidth is conserved and latency is reduced. Quality of service is also closely related to priority issues. In case that multiple service requests arrive at a service station, a priority should be assigned to each request based on properties such as emergency level, requester class, and requested quality of service.

Integration of heterogeneous information sources should be facilitated through appropriate protocols and standards of information exchange that ensure the interoperability of the underlying systems. Data, presentation, control, and functional integration issues [40] come into play when users address heterogeneous sources of information. Information integration should not require any changes to the existing data repositories and applications to avoid the costly data extraction, data conversion, or application conversion. Thus, the role of mediation services that bridge the semantic gap among non-uniform systems and services is very important [27,30,31]. Overall, an integration architecture is judged by its operational effectiveness (cost, schedule, effectiveness), user performance (response, profile of information requests), supported information integration services (task, data, and syntax semantics), breadth of information coverage, and evolution.

Extensibility of information sources and services is also critical to the evolution of healthcare networks at a regional, national, and probably transnational scale. As far as information sources are concerned, extensibility applies to the *number* and *type* of information sources in the form of information systems and data repositories, which can be included in the information space. *Information coverage* is another important aspect of extensibility and refers to the precision and accuracy of the semantic mapping among the global information model and the information model of a particular information source. Information coverage frequently necessitates an iterative process, which may be at a different stage when different user groups and tasks are concerned. Thus,

it is important that the architectural framework provides tools to assist this iterative information coverage process. Finally, *service coverage* refers to interoperability among services and information sources. In this case, middleware services enable the inclusion of new services and the binding of existing services to newly attached information systems.

Finally, *task-oriented workflows* define a variety of contexts within which the previously mentioned issues should be adequately addressed. Region-wide task-oriented workflows, and the appropriate middleware services enacting them, support the efficient handling of medical acts and processes, and the tracking of patients as they are seamlessly moving within the healthcare environment [25,35].

3.3 Digital Library Technologies

The digital library community has addressed the issues and requirements described in the previous paragraph. In fact, the healthcare information infrastructure may be viewed as a distributed digital library. A number of digital library projects deal with issues related to the high-level integration of heterogeneous information sources [1,3,4,6,62]. A digital library is popularly viewed as an electronic version of a physical library. However, the replacement of paper by electronic storage, leads to major differences including storage in digital form, direct communication, and copying from the master version. Furthermore, a digital library is only one element in the process of creating, storing, accessing, selecting, and distributing information. While the technical focus of digital library research is on the central functions of storage and access, major changes will occur, in all elements of the information generation to distribution chain, in the future systems [71].

Although there is no formal definition of a digital library, one could define its basic characteristics. The collection and management of information repositories that may reside on different computers characterize a digital library. Each repository is separately organized under a different schema or not organized at all, and a *global structure* is imposed on the information space. A *single interface* is available for searching all repositories. Queries are based on the structure of the component repositories, which are described through the use of *metadata* and, sometimes, the data themselves. Finally, a digital library takes care of *interoperability* by utilizing separately the local

tools of each repository and combining their functionality and results.

Thus, heterogeneity in its various forms is one of the basic problems addressed by a digital library [4,52]. Furthermore, a digital library should conform to a set of basic principles [2], which also apply to the healthcare domain. Namely: a) a digital library exists within a legal and social framework (i.e. domain-specific), b) digital library concepts are hindered by terminology, c) the architecture of a digital library should be separate from the content, d) a digital library should utilize metadata as handles to the actual information, e) the objects delivered to the user may be different from the stored objects (i.e. abstracted information), f) each digital library repository is responsible for the maintenance of the information it contains, g) a digital library should address the concept of collection, since users are interested in intellectual work, not digital objects.

There are several digital library projects worldwide, addressing issues such as interoperability, economic models, low bandwidth retrieval, multilingual interfaces, etc. The Stanford Digital Library Project aims at providing an infrastructure that enables interoperability among heterogeneous, autonomous digital library services, such as search and information processing. An extensible metadata architecture promotes interoperability among existing and de-facto metadata standards by including attributed proxy models, attribute model translation services, metadata information facilities for search services and local metadata for information services. These components of the metadata architecture provide, exchange, and describe metadata for information services. Metadata repositories maintain structured descriptions of both the collections where the search proxies have access and the search capabilities of the proxies. Autonomous services and collections in a distributed heterogeneous environment are accessed indirectly by proxies, which facilitate compatibility and interoperability through the utilization of metadata.

European research organizations and national libraries are leading participants in an international effort to create a standard container architecture for metadata. A recent workshop, held in April 1996 at the University of Warwick, was organized by the United Kingdom Office for Library and Information Networking (UK) and the Online Computer Library Center (USA). The participants included representa-

tives of several national libraries as well as software companies and universities. The Warwick workshop achieved consensus on a limited core of description elements (the so-called Dublin Core) and on the basic design of a container architecture for encapsulating these core elements with other sets of metadata [38,70]. Examples of such metadata, or 'data about data', include library catalogue records, specialized descriptions (e.g., for maps), terms and conditions of use, pricing and payment information, and labels for sexual or violent content. The architects of this so-called Warwick Framework are seeking to implement the container concept in a variety of ways. The simpler implementations will use existing HTML tags and compound MIME-typed documents, and will work on today's World Wide Web with minor extensions. More powerful implementations will use a CORBA-like [14] distributed object framework. A proposed Nordic Metadata Project between the national library systems of Sweden, Finland, and Norway is developing and extending this system, adapting it to Scandinavian formats and cataloguing standards.

As an example of evolvable, integrated, added-value services of educational value, consider the Visible Human Project [1,69], that has created a digital medical library including CT, MRI, and anatomical images of human cadavers. The long-term goal of the project is to make the print library and the image library a single, unified resource for health information. To achieve this goal, methods will be developed to link image data to symbolic text-based data, which is comprised of names, hierarchies, principles, and theories. Hence, a system of knowledge structures, which will transparently link visual knowledge forms to symbolic knowledge formats, will be created. However, although experiments are being done using generalized linkage methods like the use of hypermedia where words can be used to find pictures and pictures can be used as an index into relevant text, standards do not currently exist for such linkages.

There are two architectural models for provisioning networked applications, facilitated through vertical and horizontal integration [45]. In the extreme form of vertical integration, a dedicated infrastructure is used to realize each application. In contrast, one or more integrated bitways, a set of middleware services, and a diverse set of applications characterize the horizontal integration model. An important feature of horizontal integration is

open interfaces, whose purpose is to contribute to modularity through the separation or independence of the application from the “execution engine” upon which it runs, which we call “platform independence.” A general approach that has been advocated in many recent research projects dealing with networked heterogeneous information servers, examples are the TSIMMIS project at Stanford [54] and the Garlic project at IBM [9], is to identify the functions of mediators and wrappers or translators. Typically, wrappers or translators convert data from multiple sources into a common model, and mediators integrate data provided by multiple translators, in the common model [22]. The corresponding domain-specific middleware services facilitate multiple views to the underlying information sources realizing the so-called cooperative information systems, which are both flexible and efficient.

The approach currently supported by a number of standardization bodies, digital library initiatives, and R&D projects funded around the globe, model the HII as a distributed collection of objects which provide the appropriate framework for integrating resources in these environments. The target framework naturally implies a three level hierarchy: networked applications and user-oriented services, middleware services, and infrastructure or bitways [42,45]. Networked applications are distributed across a heterogeneous telecommunications and computing environment, providing value to the user through collections of functionality. Middleware is a horizontal layer residing on top of a set of networked computers, providing

a set of enabling services with standard programming interfaces and communication protocols, even though the underlying hosts and OS may be heterogeneous [17]. Middleware services provide functionality of a generic or supportive nature that is available for use in building all applications and is provided as part of a computing and telecommunications infrastructure. Examples of middleware services would be audio or video transport, file-system management, printing, electronic payment mechanisms, encryption and key distribution, and reliable data delivery. Bitways are network mechanisms for transporting bits from one location to another. In this framework, basic and emerging system integration technologies (e.g. CORBA, DCOM) will alleviate the ‘islands of technology’ syndrome; while methods and tools providing middleware and migration paths for the integration of legacy systems will lead to truly open and interoperable architectures. It should be noted however, that the placement of a particular service in one of the three layers is not strict: the advancement of technology causes services to move between layers.

Hence, existing and emerging digital library technologies provide a promising approach to the complex problems involved in the development of the healthcare information infrastructure. They have been used in the design of a reference architecture to enable the provision of integrated user-oriented telematic services in the healthcare network of Crete, which is currently under development.

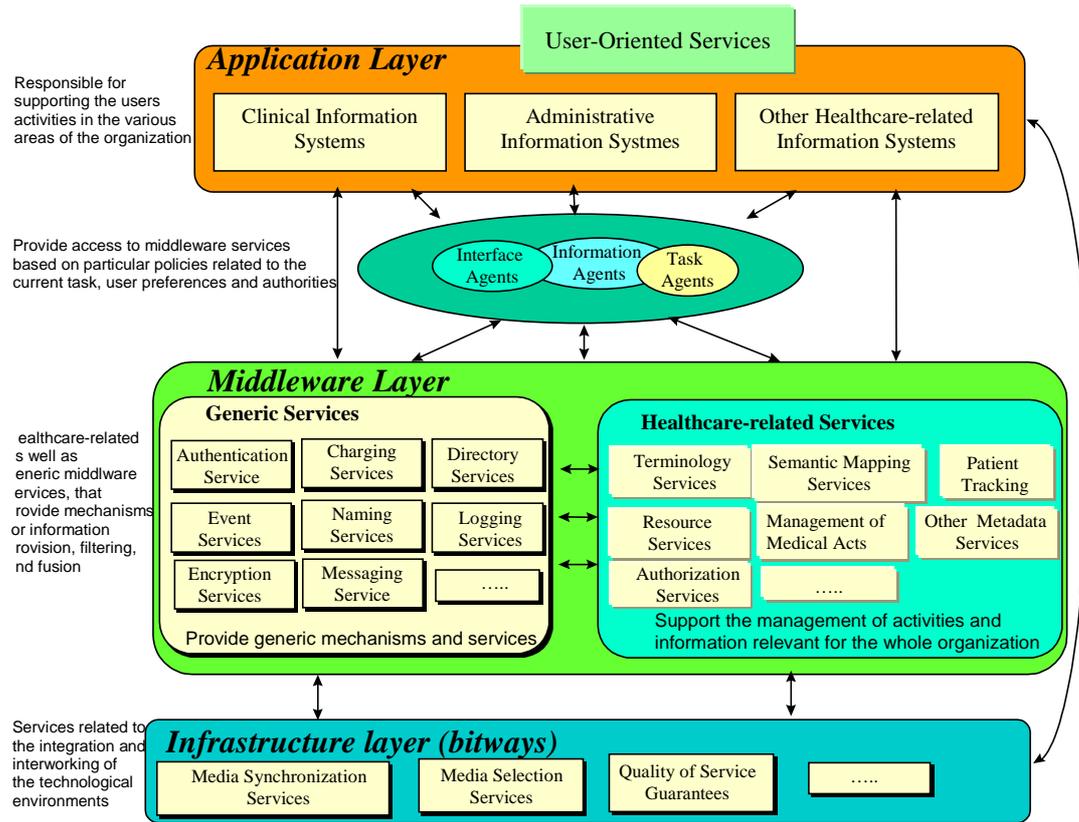


Fig. 3: The reference architecture designed for the provision of integrated user-oriented telematic services in the healthcare network of Crete.

4. The Integrated Healthcare Network of Crete

4.1 Elements of the Reference Architecture

Information technology planners use the term architecture to refer to abstract technical descriptions of systems which, instead of including the level of detail needed for construction, are conceptually-based and provide the basic framework for the creation, movement, distribution, use, and management of information [75,76]. A reference architecture model describes a system in terms of basic functional elements and the interfaces among them. It clarifies where protocols need to be defined and identifies groupings of functionality, without implying a specific physical implementation.

The reference architecture for the regional network of Crete appears in Fig. 3. The application layer of the reference architecture shows the multiple points of access to the information space, contributed by clinical and administrative information systems, user-oriented serv-

ices, etc. Besides providing points of access to the information space, some of these applications are also information sources. Specific applications employ DBMS to store various information objects, such as multi-dimensional image data, text, audio, video, etc. and appropriate interfaces for data acquisition from medical modalities that provide data and image collection services. Effective information access involves rich interactions between users and information residing in diverse locations. Users seek and retrieve information from the sources and use various tools to browse, manipulate, reuse, and generally process information. Currently, the WWW is moving to an object-oriented environment (OMG CORBA or ILU) for better services and flexibility in adding new services. The web approach provides a simple way and interface to access multimedia data and to query and browse the information sources.

Individual clinical information systems are associated with their own data model and user-interface. Dedicated applications or application components provide user-oriented services, a customizable point of access to the information space. User-oriented services may be considered as the means to facilitate secure,

task-specific information access through metadata and agent-based technology. This approach places user-oriented services one level higher than information systems, which are the sources of information. Nevertheless, in Fig. 3, user-oriented services appear in the application layer, to emphasize their role as an access point to the information space.

In the middleware layer of the reference architecture, key services support several aspects of user interaction with the information space. According to the Healthcare Information System Architecture European prestandard of CEN [CEN95], two classes of components are identified in the middleware layer of any healthcare information framework, namely Healthcare-related Common Components (HCC) which support the applications with services relating to the particular activities of the healthcare domain, and Generic Common Components (GCC), which support the applications with general-purpose facilities, common to any information system in any type of business domain. Extending the approach proposed by CEN, we are considering generic and domain-specific middleware services, which employ metadata and agent-based technology. Metadata (semantic mapping) services provide the framework for the support of meta-information related to content, provenance, form, functionality, and usage statistics [58]. Meta-information is needed not just to support integration across disparate sources and services, but just as importantly, to support a number of other activities in an information workspace including selecting, understanding, utilizing, and remembering sources and their contents. In the adopted reference architecture, generic services facilitate the use of various forms of metadata in the context of the global healthcare domain reference model, discussed in section 4.2.2. Naming and indexing/searching services are essential for the efficient retrieval of objects. Examples of such services are: Handles [38], DNS (Domain naming Service), CDS (DCE naming service), and CORBA name services [21]. An alternative would be to use global naming services (unique object identifiers) for accessing various objects. In the reference architecture, appropriate agents will help users search for objects using the metadata, naming and index services offered by the architecture.

Access to the content maintained by information systems in different authorization domains must involve secure mechanisms to browse and use data. The strict security mechanisms re-

quired in the healthcare domain necessitate the synergy of healthcare specific services such as authorization services, with generic services such as authentication and encryption services. The healthcare-specific patient tracking service employs generic directory, naming, event, and logging services to keep a consistent view of the patient record segments maintained by heterogeneous autonomous information systems region-wide. Standardized directory services that are being considered are LDAP and CORBA-based.

Terminologies, vocabularies, and thesauri are important in healthcare. There is a number of on-going standardization efforts regarding clinical vocabulary, including the HL7 initiative. Therefore, terminology-mapping services, which map one terminology standard to another, need to be established. In the context of the reference architecture, terminology services offered by terminology servers supporting evolving standards such as ICD10 and UMLS [33,48], will help healthcare professionals and researchers retrieve and integrate electronic biomedical information from a variety of sources. The UMLS project of the National Library of Medicine develops machine-readable 'Knowledge Sources' that can be used by a wide variety of application programs to overcome retrieval problems caused by differences in terminology and by the scattering of relevant information across many databases. The goal is to make it easy for users to link disparate information systems, including computer-based patient records, bibliographic databases, factual databases, and expert systems.

The infrastructure layer provides the basic technological facilities (hardware and software) enabling the proper functioning and interactions of the various components. Sample services provided by this layer relate to media synchronization, media selection, and quality of service guarantees. Because bitways support a variety of applications, each based on different standards for what the user finds acceptable, it is usually assumed that bitways provide variable QoS (a different QoS to each application) [44]. Quality of service guarantees assumes that resources are expensive and need to be conserved. Therefore, resource-allocation, pricing, and billing mechanisms are necessary to adjust resources such as bandwidth, buffer space, etc. to the specified QoS.

Generic and healthcare-specific protocols are used for the communication among the components of the architecture. Generic protocols for communication include SQL, Z39.50, HTTP,

and DIENST, while DICOM 3 [17], HL7 [28], and EDIFACT are healthcare-specific. The protocols are used to interact with various databases. Some of them are legacy databases, which have to be integrated seamlessly into the overall architecture. Furthermore, CORBA, CORBAMED, and other healthcare-related standards, provide the necessary framework for the seamless interworking of generic and healthcare-related services, thus ensuring an open and extensible framework that conforms to standards and supports evolvable and customizable user-oriented services.

4.1.1 Applying the Reference Architecture: A Possible Scenario

The reference architecture employs agents, metadata, intelligent retrieval strategies, and middleware services to provide user-oriented services. The user only needs to contact one of the task agents, which communicates (using a protocol like CORBA) with other specialized agents to get all the information needed to fulfill the user request. Specialized middleware services or agents provide directory information, information about available hospitals and patient records, doctors, expense coverage by insurance policies, existence, use and prices of drugs, etc. Most of this information, like the electronic patient record data, are not present on one server, so an agent may have to find the involved healthcare organizations and send them individual requests, using a protocol like CORBA (see Fig. 4).

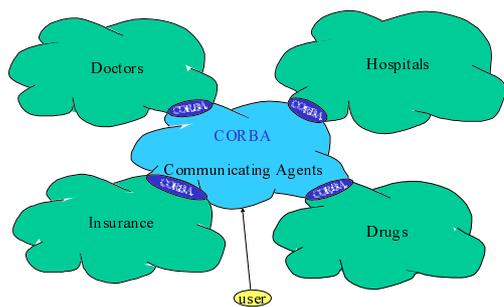


Fig. 4: An integration technology like CORBA will be utilized to integrate information sources.

To further illustrate the use of the reference architecture as the means to realize the healthcare information infrastructure, we will describe a healthcare scenario, which involves patient-data, resource, educational and collaboration services. Consider a medical practitioner located at a remote site facing a difficult case. The practitioner examines the patient and cre-

ates a patient folder in the primary healthcare information system. Although past medical history is not available, the patient recalls visiting other regional healthcare centers.

The practitioner needs to locate and view the virtual patient record of the patient. Based on the information in this folder and his current observations, relevant data from past healthcare episodes are retrieved and organized. The practitioner reviews the retrieved information, and sends e-mail to the doctors that treated some of the encounters. The information on the retrieved records is alarming, therefore, the practitioner decides to consult an expert. The practitioner through a resource service, views the list of available experts, and holds a consultation session with one of them using a user-oriented collaboration service.

Fig. 5 shows the partial execution of this scenario in the reference architecture. When the practitioner enters the visit record, the interface agent associated with the system asks whether the retrieval of similar cases might aid the diagnostic process. Upon a positive answer appropriate task agents are contacted. The task agents use terminology, metadata (HRDM), and directory services to determine the relevant information sources. Appropriate queries are placed, and the results are filtered through information agents. The interface agent presents the user with the list of similar cases. Based on the retrieved cases, the practitioner may contact via e-mail the treating physician of a case, or choose to book a teleconsultation session with a field expert.

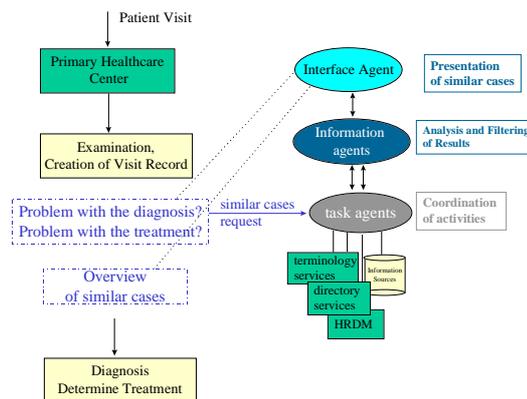


Fig. 5: User-oriented services will employ agent-based technology, metadata, middleware services, and intelligent retrieval strategies to aid medical professionals in patient management.

4.2 Evolution of the Healthcare Network

The Center for Medical Informatics and Health Telematics Applications (CMI/HTA) of the Institute of Computer Science (ICS), Foundation for Research and Technology - Hellas (FORTH), is involved in the development of an integrated healthcare network in the region of Crete (Fig. 6), which can serve as a model for national and transnational healthcare networks. Related research and development activities focus on the intelligent management of medical data in distributed multimedia databases, the real-time resource management in regional or national healthcare networks, on methodologies and architectures for the integration of heterogeneous information systems, the processing and analysis of multimedia medical data, particularly 2D and 3D images, and the indexing and retrieval of medical images based on their content. Additional research effort is focused on developing explicit models of the underlying workflows in the various environments (hospitals, regional healthcare networks, Internet) required for the optimization of the process of information distribution, thus improving the availability and accessibility of patient data.

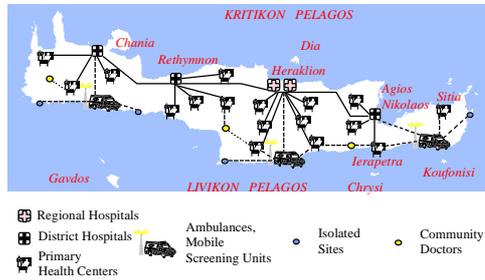


Fig. 6: Topology of the Regional Healthcare Network of Crete.

In this context, a number of autonomous applications have been developed to address specific needs of the healthcare domain. These applications include: a) clinical information systems, b) cooperative working environments, and c) added-value services. Individual systems were initially developed as standalone applications with minimal attention given to integration, reusability, and distribution issues. In the current context of the integrated healthcare network of Crete, which is under development, these systems are gradually integrated in the proposed reference architecture utilizing common components and services. Hence, the issue of integration and interoperability of heterogeneous autonomous distributed systems (HADS) is addressed in a generalized framework de-

veloped by digital library technologies. In this way, a regional digital library providing user-oriented services using metadata, middleware services, and agent-based technology, will evolve out of heterogeneous autonomous distributed systems.

In the following sections, we overview the bitway and middleware services currently under development, we describe our approach to the adoption of agent-based technology, and outline the basic features of some of the applications and enabling services currently under various stages of development in our Center.

4.2.1 Bitway services

Bitway services provide the capability for distributed searching and retrieval based on QoS (quality of service) as defined by users, as well as mechanisms for efficient usage of resources such as processing power, I/O and memory in the servers, and bandwidth and network buffers [19,60]. Important considerations in designing novel performance-based architectures include response time in obtaining the objects of interest, and quality in viewing multimedia objects. This implies a careful study of performance issues and dynamic routing of user requests to obtain the relevant documents. The architecture has to support mechanisms for adaptive (based on demand) replication of meta-information, indexing information, etc. Replication also includes mechanisms for caching.

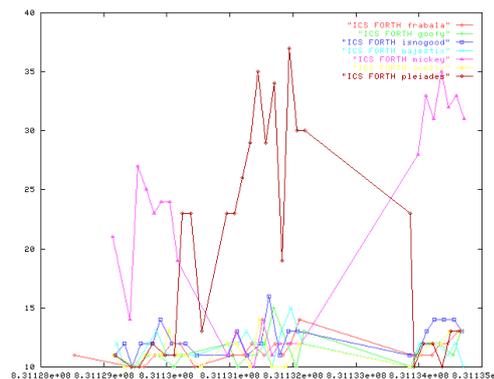


Fig. 7: Number of requests versus elapsed time to each of the 7 DIENST servers. The hosts on which the DIENST servers reside are shown (ICS-FORTH.hostname).

Currently, we are using a testbed based on the DIENST system services for research and development in the following areas: performance management, load balancing of requests, scalable indexing architectures and QoS in searching and presentation. A testbed of 7

DIENST servers is being used for experiments. Fig. 7 shows the load (requests per hour) on the DIENST servers over half a day. Such performance metrics will be used for load balancing user requests in the healthcare network of Crete.

4.2.2 Middleware services

In the middleware layer of the reference architecture, several vital services are at various stages of development. These are briefly described in the sections below.

4.2.2.1 Healthcare Domain Reference Data Model (HRDM)

Integration of information and knowledge from different sources is becoming increasingly important and remains one of the most intensely contested information technology issues in the international research arena [4,54,61]. A promising approach to the problem of harnessing disparate information sources is to employ meta-level descriptions of the information sources. As proposed in [30,31], gaining control of an organization's information resources at a meta-level allows autonomy of individual systems at the data-instance level. The objective of the meta-model is to achieve enterprise information integration over distributed and potentially heterogeneous systems, while allowing these systems to operate independently and concurrently. Hence, semantic interoperability is ensured, since the requester and the provider have a common understanding of the requested services and data.

In the regional network of Crete, the semantic mapping of the diverse information sources is based on the Healthcare Domain Reference Data Model (HRDM), a research prototype under development in our laboratory [51]. The HRDM provides an extensible meta-model of the information that is maintained by the participating HADs. This information facilitates semantic mapping of the component information systems through a global data model. Thus, the relevant data may be configured differently at different locations and can be semantically mapped onto a common model, before they are presented to the user. As the information model of new information sources is introduced, the HRDM is further expanded.

Given the proliferation of terminologies in healthcare, the selection of a single terminology globally applicable in all clinical domains is an important problem. It is exactly this

problem that we are addressing with HRDM. We allow each individual information source (i.e. information system) to be optimized for the specific needs of a single functional unit, including adherence to the most appropriate terminology standards. For example, it would today be very hard to envisage a PACS system not compliant to DICOM. In addition, by providing mechanisms (mapping services and mediator agents) for expressing the schemata of individual information systems into the HRDM, the latter will gradually evolve into a meta-thesaurus maintaining the semantic interrelationships between various terminologies and vocabulary standards. The dynamic extensibility provided by the SIS (see later in this section), on which the HRDM is implemented, allows for additional information sources (and therefore terminologies) to be expressed in the HRDM. Therefore, semantic interoperability between the relevant terminologies employed by the various autonomous information sources is achieved.

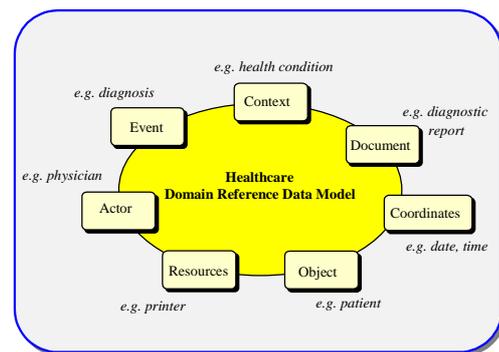


Fig. 8: The basic entities of the Healthcare Domain Reference Model (HRDM).

As shown in Fig 8, the basic entities of the HRDM are:

- *Actors* perform actions relevant to the provision of healthcare. An actor may be a person, a group of persons, or an institution. An actor may be located anywhere and be responsible for or own a number of resources. Typical examples are: a physician, a nurse, a hospital, an insurance company, etc. Note that an actor is not the recipient of medical care.
- *Object* is the recipient of healthcare. An object may be located at a specific site or be assigned to an *actor*. Typical examples are: a patient, a population, etc.
- *Event* is a series or a group of incidents that change the state of the system. An event happens at specific *time* and *place*

coordinates, usually affects an *object*, and may be handled by an *actor* or use a *resource*. Typical examples include medical acts, diagnosis, patient admission, document update, etc.

- *Context* is anything that occurs at certain *time* and *place coordinates* and does not affect the state of the system. Typical examples include a health condition, a disease, a symptom, an activity, etc.
- *Resources* include materials that are used in the process of providing medical care. A resource may be associated with a *place coordinate* or be applicable to an *object*. Example resources include medical equipment, printers, drugs, and experts.
- *Documents* include information about or results of medical procedures. A document may also include information about an *actor* (a person or institution) that is involved in the process of healthcare provision. It may be in paper or electronic form or even in the form of raw data as produced directly from medical equipment or after some processing. Typical examples are an image, a diagnostic report, an insurance contract, etc.
- *Coordinates* specify place and time. They may refer to an *actor*, an *object*, an *event*, a *resource*, a *context* or a *document*.

Important categories of metadata that should be described in the framework of the HRDM are content, provenance, form, functionality, and usage statistics [18]. *Content* refers to the information that is covered by the source, since the breadth of information coverage may not be the same for all information sources. *Provenance* refers to the nature of the processes that produced and/or maintain the various information sources. *Form* is relevant to the schema for items contained in the source, including the available attributes and the types of values stored for those attributes. Note that, for the same source, it is possible that multiple metadata views to the underlying schema coexist to facilitate alternative views to the information space. *Functionality* refers to the supported information integration services and medical acts, with emphasis on the capabilities and properties of the accessed services including the kinds of the searches supported and performance characteristics. Finally, *usage statistics* provide information about source usage including previous use by the same user or other users. These statistics are often more useful if aggregated by categories of users or by organizational or other hierarchical groupings. The use of specific resources and services may

be also logged, to permit the drawing of conclusions on the behavior of particular users or user groups in the course of a medical act.

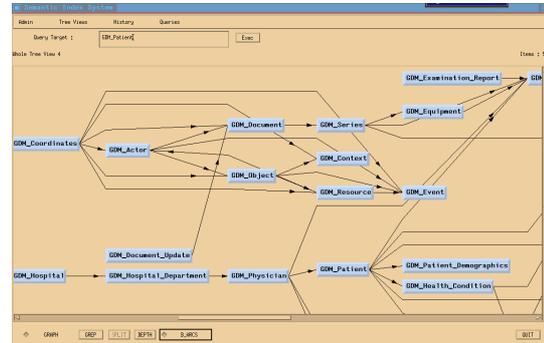


Fig. 9: The X-Windows user-interface of the SIS, showing parts of the HRDM.

The HRDM has been developed using the Semantic Indexing System (SIS) [13]. The SIS is a tool for describing and documenting large evolving varieties of highly interrelated data, concepts, and complex relationships. It consists of a persistent storage mechanism based on an object-oriented semantic network data model, and a generic interactive user interface to insert and retrieve information in various ways. The SIS offers significantly richer referencing mechanisms than relational or ordinary object-oriented systems. Together with the very high query speed along references, these mechanisms allow the data and schema to be kept redundant-free. One of the noteworthy features of the SIS is its uniform treatment of data and schema by the data entry, the query system and the user interface.

The SIS data model is characterized by an unbounded instantiation hierarchy, multiple strict inheritance (generalization), multiple classification, and multivalued attributes which may have their own attributes. The generic user interface (see Fig. 9) supports menu-guided and forms-based query formulation with graphical and textual presentation of the answer sets, as well as graphical browsing in a hypertext manner. Menu items, menu layout, and domain-specific queries are user-configurable. A forms-based interactive data entry facility allows for entering data and schema information in a uniform manner. SIS exploits widely self-adaptation to the schema or meta-schema and it is customizable to application specific tasks. Thus, applications can be built without code writing. For heterogeneous data-exchange, a configurable SQL-based tool transforms views from RDMS into SIS repre-

sentations and imports and updates these items in an SIS database. The current version of the SIS has been tested with a population of up to 5 million entities and references. Up to a total population of 500000 references, no significant influence of the population size on the query speed could be measured.

4.2.2.2 Management of Medical Acts

Workflows are activities involving the coordinated execution of multiple tasks performed by different processing entities [25]. In the healthcare domain, activities correspond to worklists, tasks correspond to diagnostic service requests, and processing entities are the various autonomous information systems. Important requirements for the efficient and reliable operation of applications supporting workflow management include deep understanding of the processes to model as well as workflow implementation and automation. The separation of work activities into well-defined tasks, rules, and procedures allows for the modeling of healthcare processes in a rigorous and comprehensive way. This is a prerequisite for a totally automated information system [67].

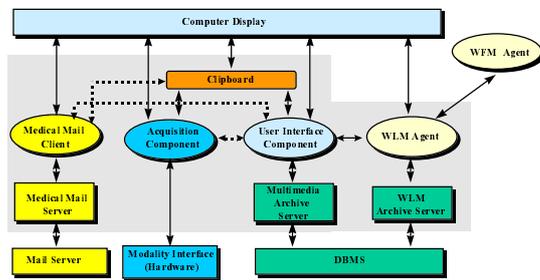


Fig. 10: Workflow management in an integrated hospital environment.

In this context, a CORBA compliant Workflow Management (WFM) agent that provides a domain-specific middleware service for the management of medical acts has been designed and implemented. The WFM agent, shown in Fig. 10, operates as a coordinator of Work List Management (WLM) agents distributed hospital-wide [36]. The WFM agent maintains information on the services provided in the integrated hospital environment and acts as a gateway among heterogeneous information systems. This domain-specific middleware service enables clinical practitioners to know what services are offered, in which department, during what times, and who is responsible for their execution. In addition, the WFM supports

a set of services capable of providing distribution of available services in the integrated hospital environment, worklist decomposition and routing, appointment reservation, as well as report composition. The WFM also keeps a record of all the intra-hospital examination request and report traffic. The WLM agents cooperate with the WFM agent to coordinate the execution of hospital-wide processes, such as patient admission, placement and management of diagnostic service requests to the appropriate clinical information systems, etc. A diagnostic service request can be either a specific medical examination or a set of medical examinations (study). After the completion of the examinations, which correspond to the steps that compose the requested procedure, the final report is constructed and transmitted to all its recipients. The diagnostic service request is considered complete only after all recipients receive the report. The CEN/TC251/WG4's MEDICOM standard has been employed for medical imaging communications, and the Health Level 7 (HL7) standard has been used for character-based, application-level message exchanges.

4.2.2.3 Regional Healthcare Resources

A global resource database is being developed as a domain-specific middleware service. It maintains information on regional healthcare resources used in medical activities. Such information includes physical resources, such as hospital departments, locations, medical modalities, mobile units and their characteristics, etc., and human resources, such as healthcare practitioners, including their roles, data usage, and rights. It constitutes a vital middleware service used by a number of applications or other services, such as a Prehospital Health Emergency Management System, a CSCW environment, an asynchronous teleconsultation component, etc., described in the sections that follow.

4.2.2.4 Authorization and Security services

The definition and control of the authorities of individual users in the execution of various medical acts as well as in accessing parts of the patient record, represents a major concern in the healthcare environment [57,73]. In fact, besides the need for supporting the diversification of roles and responsibilities which is typical of any large organization, additional ethical and legal aspects can be identified in the healthcare domain due to the nature of the managed information. On the basis of such

considerations, two fundamental and complementary needs can be identified for any healthcare information system: a) the security of the managed data and, b) the control and monitoring of the actual authorizations for the individual users executing certain tasks on the system.

The authorization, authentication, and encryption middleware services of the architecture aim at supporting this specific need by providing:

- a comprehensive and consistent repository where an authorized individual defines the rules according to which the diverse user groups and individual users may execute the various functions provided by a system (domain-specific authorization service).
- a standard mechanism according to which data is managed by the overall system to ensure an adequate level of reliability and protection (generic authentication and security services).

4.2.3 Agent-based technology

Information is becoming increasingly difficult for a person or machine to collect, filter, evaluate, and use in problem solving. As a result, the problem of locating information sources, accessing, filtering, and integrating information to support decision making, as well as coordinating information retrieval and problem-solving efforts has become a very critical task. In the context of digital libraries, agents have been widely employed as a means of providing assistance to users in the complex tasks of information tracking, filtering, and fusion [6,39,59]. In the reference architecture of Fig. 3, agents are differentiated in terms of their functionality and communications needs. Following the recommendation of [63], three types of agents have been defined: interface agents, task agents, and information agents. *Interface agents* interact with the user, receiving user specifications and delivering results. They acquire, model, and utilize user preferences to guide system coordination in support of the user's tasks. *Task agents* help users perform a task by formulating problem solving plans and carrying out these plans through querying and exchanging information with other agents. *Information agents* provide intelligent access to a heterogeneous collection of information sources.

As already mentioned, the execution of general healthcare processes involves the cooperation of WFM agent with dedicated task agents

which collect relevant information from the user, in order to initiate, monitor, and control the execution of a task.

4.2.3.1 Next Generation Medical Workstation (NG -MW)

It is often necessary to view the integrated medical record of a patient in a variety of different ways, e.g. encounter-centered, problem-centered, task-centered, patient-centered, and as temporal snapshots. The use of anticipatory computing in the design of the healthcare professional workstation has been introduced in a hallmark 1995 paper by Blois [26]. In this paper, Blois challenged informaticians to devise and develop a workstation that could predict user needs and provide an environment that knows about its users, their preferences and what information they will need within a given context [26].

Such a workstation is only beginning to be realized. The NG-MW should use presentation strategies, models of communication goals, knowledge of the user's current schedule and context, and available resources to optimize *information transfer* to a user. Specialized interface agents should support access to *context-relevant views* of the entire patient record, regardless of whether the data were retrieved from one database or a number of networked information systems. The development of such a medical workstation, incorporating appropriate agent-based technologies, is an important objective of our future work.

4.2.4 Clinical Information Systems

As far as the application layer of the architecture is concerned, several clinical information systems have been developed and utilized as autonomous networked applications. Appropriate extensions or modifications to these information systems are currently carried out, so that they are seamlessly integrated to the proposed reference architecture. The degree of conformance to the reference architecture indicates the capability of the information system to be integrated, from both the functional and information viewpoint, with the rest of the information systems. System conformance is measured by the capability of an information system to exploit, when necessary, one or more of the common components specified in the reference architecture. In the following paragraphs, the basic characteristics and functionality of these information systems, currently in the process of an incremental evolution to-

wards this reference architecture, are presented.



Fig. 11: Typical screen from IMACS.

4.2.4.1 Image Management and Communication System (IMACS)

IMACS is an integrated system for image acquisition, management, processing, and communication. Using the DICOM 3 standard, IMACS can acquire images from different diagnostic modalities. The system uses a multimedia database management system to organize and manage data. Data is hierarchically structured, by grouping patient-related multimedia data into objects such as exams, studies, and acquired image sets, according to the DICOM 3 standard. Its Distributed Hierarchical Storage Management (DHSM) [67] module supports intelligent storage management, based on predefined models (*model-driven*) or in response to events (*event-driven*). In the model-driven image management case, image data are automatically moved downward the storage hierarchy according to a Least Recently Used (LRU) algorithm. By contrast, in the event-driven image management case, the DHSM prefetching algorithm identifies the related patient data managed by the IMACS system and selectively transfers them to the appropriate server, so that they are available during reporting together with image data generated by the current imaging procedure. In addition to its use in a distributed hospital environment for the management and communication of patient-related information, the IMACS system is also used as a Teleradiology/Telemedicine system for inter-hospital communication of multimedia medical data and remote consultation. Specific common components and middleware services are currently being introduced to IMACS. Hence, it is being converted from a vertically integrated applica-

tion into a horizontally integrated application within the reference architecture.

4.2.4.2 Laboratory Information System (LIS)

The Laboratory Information System (LIS) is an autonomous application responsible for the acquisition, management, and communication of laboratory data. The IEEE MEDIX data model was adopted in the design of this application. Data acquisition is performed through the standard RS-232C serial interface provided by most laboratory instruments. After the necessary two-phase validation procedure, for eliminating transmission and data errors, data are transferred into the patient record. LIS eliminates the need for data re-entry in examination scheduling by employing domain-specific services for the management of acts. Thus, as far as functional integration is concerned, LIS is compliant with the reference architecture. Information integration will be achieved by performing the semantic mapping of its data model into the HRDM.

4.2.4.3 Intensive Care Unit Information System (ICU IS)

The ICU information system, currently in its later stages of implementation, is a CORBA-compliant information system consisting of three autonomous components. Each of these components is responsible for managing different data types. An integrated bedside display for the practitioner is thus provided for viewing patient demographic data, and vital signs of both static measurements and continuous recordings. Support is provided for several viewing modes: real-time data display, daily overview, as well as for charting mode [36]. Functional integration of the system with other autonomous systems (IMACS, LIS) is achieved by using the management of acts domain-specific middleware service.

4.2.4.4 Prehospital Health Management System

Even though the Prehospital Health Emergency Management System (PHEMS) [41] has been designed as an autonomous system, it operates in the context of the regional healthcare network of Crete, where integration and interoperability issues are of major importance. In addition to the effective resource management and efficient response to emergency service requests, the PHEMS provides medical services for access to clinical multimedia data

stored in patient records, for the acquisition of vital signs and the selection and application of predefined medical procedures. By means of these services, clinical multimedia data from the patient record archive, as well as currently acquired vital signs and real-time pictures from the scene of an accident, can be exchanged between the healthcare emergency unit, the hospitals, the primary care centers, and the mobile units. Locally stored predefined medical procedures are triggered automatically by conditions resulting from the analysis of vital signs at the assistance site or remotely by a doctor at the Health Emergency Coordination Center (HECC).

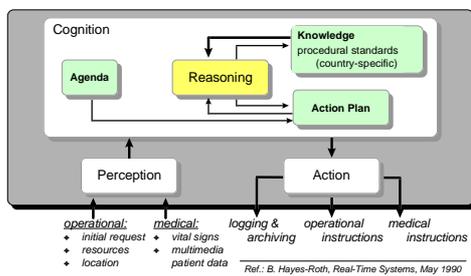


Fig. 12: The perception-action paradigm is utilized in prehospital healthcare emergency management.

The architecture of the system uses a perception-action paradigm illustrated in Fig. 12. Perceived operational data may be incoming calls, i.e. emergency service requests, but also information about available resources or data describing the current location of mobile units, patients, etc. Perceived medical data include the acquisition of vital signs, biosignals and patient image data, and retrieved clinical multimedia data. The cognition component of the architecture gives instructions and recommendations based on the perceived information and knowledge about operational and medical procedures stored in ruled-based systems and more sophisticated knowledge bases. Typical problems to be solved by the cognition component are the efficient and effective management of resources under dynamically changing conditions or the determination of the adequate set of therapeutic or life-saving actions, triggered by continuously acquired vital signs. Finally, operational actions include, for instance, the allocation of resources such as mobile units, the exchange of operational instructions between the HECC and the mobile units, and the maintenance of the emergency episode archive.

The PHEMS is being developed in compliance with the reference architecture and employs

agent-based technology and domain-specific middleware services for authority assessment, resource availability and allocation, patient-tracking, etc.

4.2.4.5 Primary Healthcare Center Information System

Primary Healthcare Centers (PHC) play an important role in the provision of healthcare services within integrated regional healthcare networks. The CMI/HTA in collaboration with the Department of Social Medicine of the University Hospital of Heraklion and users from several primary healthcare centers on the island, is developing an integrated PHC information system for collecting, managing, analyzing, and communicating healthcare information generated at a PHC. The system, which manages a segment of a patient's record, is being developed with adherence to relevant standards, e.g. ICD-9 and ICD-10 for classification of diseases, DICOM 3 for image management, SCP-ECG for ECG management and communication.

Added-value user-oriented services are also being developed, utilizing the global data produced at the primary healthcare level, to allow for the assessment of the health status of the population and the early identification of healthcare problems. Furthermore, generic middleware services for messaging facilitate the communication of patient data to medical assistance companies and other healthcare actors. The PHC information system employs domain-specific terminology, authorization, patient tracking, and resource services, and domain-independent messaging, authentication, and encryption services. The uniform view of the patient record will be accomplished through the domain-specific semantic mapping service.

4.2.5 User-Oriented Collaboration Services

4.2.5.1 Asynchronous Teleconsultation

As part of our effort to develop generic and reusable tools, a domain-specific multimedia communication component, supporting remote opinion request, has been developed. As a result asynchronous teleconsultation is facilitated through generic and domain-specific middleware services, which facilitate the exchange of multimedia medical data as Messaging Application Programming Interface (MIME) attachments over the Simple Mail Transfer Protocol

(SMTP), utilizing generic encryption and authentication mechanisms.

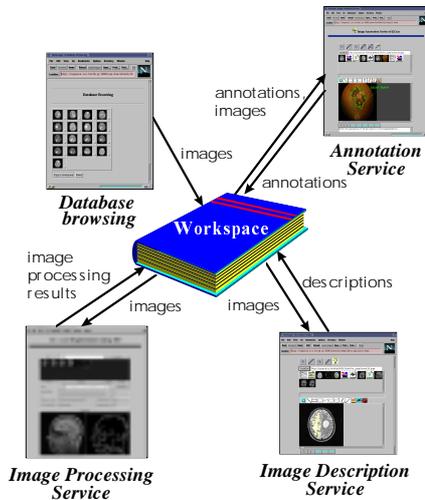


Fig. 13: Workspaces facilitate the integration of web-based medical services.

4.2.5.2 Web-based Collaboration Services

In the regional healthcare network of Crete, autonomous image repositories are gradually evolving, as users submit interesting medical images and related patient data. In this context, user-oriented services supporting collaboration are provided to authorized users through an unmodified web browser. Web-based services are integrated in the sense that both asynchronous collaboration through e-mail, postings, and shared workspaces, as well as synchronous collaboration through an on-line collaboration forum are facilitated (see Fig. 13). Furthermore, the presence of an Internet gateway leads to other regional networks and the Internet. The gradual integration of domain-specific middleware services will facilitate the discussion of medical cases and promote remote opinion request and teleconsultation in a flexible environment [11].

Shared Workspaces. Shared workspaces support the asynchronous cooperation of authorized users and service integration by providing secure and transparent access to a heterogeneous data collection, anytime and anywhere browser software and a network connection are available. This data collection includes service results and multimedia data objects (images, voice, video, patient exams, annotations, etc.) that were inserted by the users sharing the workspace. Thus, multiple users may collaborate over a workspace, sharing material of common interest. Workspaces persist between

user sessions, among services, and among users. All users connected to a specific workspace are notified of workspace updates as they occur and may inspect the property sheet of a workspace object that includes its author, the date it was last updated, and a brief description.

Annotation Service. The objective of the annotation service is to provide healthcare professionals with the ability not only to access medical image collections, but also to interact with imagery, creating, viewing, and communicating notes on groups of images and related patient data. The purpose of an annotation is to provide a rich way for users to express a comment, a refutation, a correction, or a confirmation through a generic annotation format (see Fig. 14). An annotation may include graphics and text grouped into multiple overlays. Furthermore, an annotation may be associated with various media types such as images, text, sound, hypertext, and video. The annotation service allows users to collaborate over groups of diagnostic images and related patient data. Hence, medical specialists may create annotations for their private collections, discuss them using e-mail, or use them in discussion forums and on-line collaboration sessions. Furthermore, the administrator of a site may browse annotation postings, select interesting ones, and link them to selected medical images.

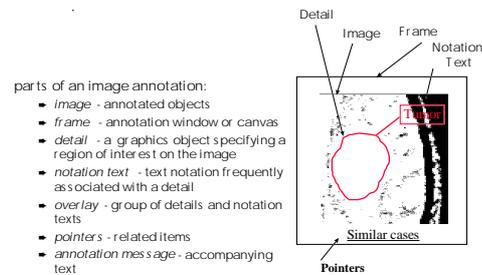


Fig. 14: The generality of the annotation format enables rich user-interaction with visual information.

On-line Collaboration. The on-line collaboration service satisfies the need for a more direct form of communication. Users share a workspace and use an “on-line talk” facility to discuss various workspace objects. In the course of an on-line session, the joint annotation of a group of images may be performed. A new instance of the collaboration server is launched, when an authorized user requests the creation of a collaboration session bound to a specific workspace. The contents of the workspace at that time constitute the conference material, and the user that requests the launch-

ing of the session controls the floor. All exchange of information passes through the server, enabling the recording of the complete session. Furthermore, snapshots of the session may be stored in the annotation format.

4.2.5.3 CoMed: Synchronous Collaboration in Medicine

CoMed [77] is based on a distributed agent architecture that enhances the performance of the shared window environment, offering short response time and reducing network load. The combination of data replication and a distributed agent architecture enables CoMed to make efficient use of the network resources during the conference, by transmitting only the events that are actually needed to maintain consistency among the various instances of the shared workspace. As a result, CoMed operates successfully in heterogeneous networks where low-bandwidth networks may be connected to high-speed networks (e.g. combinations of Ethernet, ISDN, FDDI, ATM, etc). Unlike other systems that require the group of participants to be known at the beginning of the conference, CoMed supports the addition of new members into the conference at any time.

A global authentication service supports public and private (at various levels of privacy) tele-consultation sessions, while image data integrity is ensured by the use of standard digital signature technology. Functional features address specific requirements of the medical domain, such as the sensitive, unpredictable and often urgent nature of medical consultations. In addition to mechanisms that are responsible for the management of the conference, CoMed addresses pre-conference and post-conference management issues. These include conference scheduling and data prefetching, maintaining conference proceedings, and broadcasting final conclusions to participants and other authorized personnel after the end of the conference. Through the gradual integration of middleware services for resource, authentication, and authorization, CoMed is evolving from a standalone CSCW application into a domain-specific common component that provides functionality to the user not only as a user-oriented collaboration service, but also as an synchronous medical collaboration component of the various clinical information systems.

4.2.6 User-oriented Added-value Services

4.2.6.1 DIPE: Image Analysis and Description Services

DIPE [78] is being developed to operate within an integrated healthcare network as a domain-specific added-value service, providing access to high-performance computing facilities in order to support computationally intensive image analysis and visualization tasks. The main features of DIPE include: computational resource management and intelligent execution scheduling within a local or regional network; intelligent and customizable mechanisms for the description, management, and goal-oriented retrieval of image processing software modules, so as to ease the work of medical personnel; mechanisms for the "plug-and-play" integration of already existing heterogeneous software modules; easy access and user transparency in terms of software, hardware, and network technologies; and charging mechanisms based on quality of service.

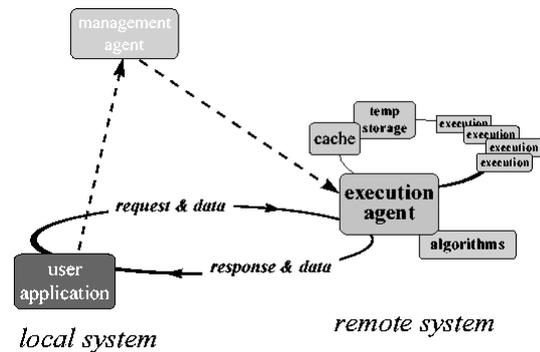


Fig. 15: The distributed autonomous cooperating architecture of DIPE.

DIPE is based on a distributed, autonomous, cooperating agent architecture (Fig. 15). It is designed so that it is modular, scalable and extensible, and it can be readily implemented on different hardware and software platforms, and over heterogeneous networks. The environment consists of a functional core, which ensures the persistent distributed execution of image processing algorithms, and can be extended to support other added-value services such as macros, resource management, algorithm retrieval, charging, etc. The core of the system consists of several communicating components: user applications, execution agents, pools of IP algorithms, and management agents.

4.2.6.2 *I²Cnet*: Intelligent image content browsing strategies

I²Cnet (Image Indexing by Content network) [50] is an on-going effort whose goal is to provide network-transparent content-based access to medical image archives, as an Internet/Intranet added-value service. Through a standard Web browser, healthcare professionals will be able to interact with image collections, browse images similar to a query image, compare these images to images from other collections, and contribute their own images or comments. Specific *I²Cnet* services currently available on the world-wide Web (WWW) include image processing and feature extraction, content- and annotation-based search for images and image-related information, and authoring of annotations and image descriptions. *I²Cnet* is organized as a network of *I²C* servers and brokers. Each *I²C* server distributes its own set of services to the brokers, which are responsible for updating the *I²Cnet* service directory. *I²C* brokers maintain directory information and service profiles based on user-feedback and access patterns and use this information to handle network-transparent service requests.

The *I²Cnet* description service enables users to import an image and specify regions of interest on it. These regions of interest may be drawn manually or can be the output of an image segmentation algorithm, which has been imported from a workspace. A number of content-extraction tools are available to the image description service. Their objective is to help the user improve the presentation of an image, view its properties, and analyze its visual content. Using the *I²Cnet* description service, users are able to create a description of an image and use it as a visual query in the *I²Cnet* retrieval by content service, to extracting similar images for the *I²Cnet* databases. The *I²Cnet* retrieval by content service allows users to retrieve from a specific image class, images with a visual information content similar to that of a visual query. Various forms of visual queries can be supported in *I²Cnet* through description types that implement specific retrieval by content strategies. A particular description type currently available on the Web provides a content-based query by example. The user interface of a retrieval by content algorithm is constructed dynamically by the *algorithm execution tool*. The result of a retrieval by content service request is a query object which includes the executed query and its result. This query object may be used to cache query results and, thus, avoid re-execution of the query if the database

has not been modified since the last execution. Furthermore, query objects may include user-feedback provided in response to query results. Hence, the system may gain information on the perspective of each user, as well as general information on how each retrieval algorithm works.

4.2.7 Discussion

The healthcare network of Crete provides a framework for the development of a HII. In this framework, information systems are sources of information, and tying them together in an execution architecture is the challenge we are currently faced with. The ultimate objective of our effort is the provision of integrated user-oriented telematic services and uniform access to networked information sources. These user-oriented telematic services in addition to providing customizable functionality to the user, will serve as common components of specific clinical information systems. To address these requirements, a reference architecture has been designed which will allow for the seamless integration of information. Digital library technologies are employed in this framework to provide the mechanisms for combining these existing systems and services into a uniform information space.

As stated in [53], there are five major dimensions in information management. These are:

- *management*: management and organization of information sources
- *retrieval*: information retrieval
- *discovery*: query formulation and refinement
- *interpretation*: semantic interpretation and visualization of queries and results
- *sharing*: information sharing and collaboration.

Toward our goal of developing the HII for the healthcare network of Crete, each of these dimensions is investigated using existing and emerging digital library technologies. The ultimate objective is to facilitate the delivery of integrated user-oriented telematic services, through the meta-level description of the information sources and the mediation of appropriate agents.

5. Conclusions

The next generation of regional healthcare networks will provide integrated user-oriented services based on autonomous and distributed information systems, knowledge-intensive ap-

plications, and large quantities of intelligently managed multimedia medical data. A key challenge facing system researchers and developers is to provide an architectural framework for the emerging healthcare information infrastructure, which would facilitate the integration of existing distributed information resources into what would appear to the users as a uniform and logical collection of information sources, supporting diverse medical applications and the provision of integrated user-oriented telematic services. 'The introduction of such services into routine clinical practice is a complex issue, whose successful resolution will require the close cooperation of technology providers and health care professionals. Then, patients will hopefully also experience the benefits of the above technological advances' [50].

In addition to addressing architectural issues, the realization of the Healthcare Information Infrastructure requires that the networking infrastructure, which consists of wide-area communications networks, wireless data networks, and mobile communications networks, will be further developed and integrated. Such an integrated networking infrastructure should provide a high degree of protocol transparency, so that both existing and emerging protocols can be accommodated. Furthermore, clinically-relevant user-oriented services provided over the HII require that information sources, rich in content, be made available and that they are properly maintained and updated by autonomous healthcare information systems. Achieving these goals will largely determine the extent to which medical expertise and clinically relevant information will actually become a shared resource, wherever it may exist. When this occurs, the healthcare sector will have become an important domain in which the real impact of the emerging global information society will be assessed. It is our belief, that in this setting, digital libraries will play a major role as enabling technologies.

In this paper, we have presented the current status of our work toward the development of an integrated reference architecture for the provision of user-oriented telematic services in the region of Crete. In the development of the regional healthcare network of Crete, emphasis is given to the creation of a framework for integrated service provision using digital libraries as the enabling technology. Advanced applications and software platforms, designed to accommodate heterogeneous processing environments and evolving user needs, are cur-

rently at various stages of development and have been described only as examples. Different service types are also being designed and implemented. These include: a) *services for dynamic content update*, i.e. services that enable the seamless creation of an integrated virtual patient record, b) *collaboration services* (asynchronous as well as synchronous) that provide teleconsultation and computer supported cooperative work between healthcare providers at different locations, and c) *added-value user-oriented services* supporting the creation of innovative content (web-based image annotation service), the processing of medical images (DIPE), and the visualization of information (NG-MW). The latter is of particular importance in the medical domain, since the medical record has to support a wide range of users with different viewpoints. Therefore, the information contained in the electronic medical record of a patient must be viewed in a variety of different ways, e.g. encounter-centered, problem-centered, task-centered, patient-centered, and using temporal snapshots.

In conclusion, the development of a HII will provide the environment for transforming unprocessed, unintegrated data into abstracted, and therefore much more useful, information. Our challenge for the future is to provide healthcare networks with integrated content, to manage this content and to make it easily accessible to various user groups. Finally, healthcare networks must be extensible and must be allowed to evolve in terms of supported functionality, services provided, and universal accessibility. An additional challenge is not only to keep up with these developments, but also 'to be creative in using this emerging information technology and telecommunications infrastructure to provide clinically significant and cost-effective added-value telematic services to the health care community, while ensuring that the potential benefit to be derived from technological advances also finds its way to the scene of an accident or the home of patients and the elderly' [ORPH96b].

Acknowledgments

The development of the Integrated Healthcare Network of Crete is a long-term goal of the Center of the CMI/HTA at ICS-FORTH. Part of the work reported in this article represents a collaborative effort also involving the Divisions of Software Technology and Information Systems, and Parallel and Distributed Systems (Pleiades) at ICS-FORTH. The authors would like to explicitly acknowledge significant con-

tributions by M. Blazandonakis, G. Doulerakis, E. Kaldoudi, D. Katehakis, E. Leisch, F. Logothetidis, A. Petrou, N. Stathiakis, X. Zabulis, and M. Zikos, all members of the CMI/HTA. Finally, this work is being supported in part by the HECTOR and ET-ASSIST projects of the European Union Health Telematics Applications Programme, and by project IHIS funded by the General Secretariat for Research and Technology of the Greek Ministry of Development.

References

1. M.S. Ackerman and R.T. Fielding. Collection Maintenance in the Digital Library. In Proceedings of Digital Libraries '95, 1995.
2. W.Y. Arms, C. Blanchi, and E. Overly. An Architecture for Information in Digital Libraries. D-Lib Magazine, February 1997.
3. M.Q. Wang Baldonado, C.-C.K. Chang, L. Gravano, and A. Paepcke. The Stanford Digital Library Metadata Architecture. Technical report, Computer Science Department, Stanford University, January 1997.
4. M.Q. Wang Baldonado and S.B. Cousins. Addressing Heterogeneity in the Networked Information Environment. Technical report, Computer Science Department, Stanford University, December 1996.
5. M.Q. Wang Baldonado and T. Winograd. SenseMaker: An Information-Exploration Interface supporting the Contextual Evolution of a User's Interests. Technical report, Computer Science Department, Stanford University, December 1996.
6. W. P. Birmingham. An Agent-Based Architecture for Digital Libraries. D-Lib Magazine, July 1995
7. A. Bernstein. Middleware: A Model for Distributed Systems Services. Communications of the ACM, 39(2):86-98, February 1996.
8. M. Blazantonakis, V. Moustakis, and G. Charisis. Seamless Care in the Health Region of Crete: The STAR* Case Study. In Proceedings of MIE'97, Porto Carras, May 1997.
9. M. Carey. Garlic: A Middleware Multimedia Information System. Presentation Slides, IBM Almaden Research Center, Oct. 1996.
10. Medical Informatics - Healthcare Information System Architecture, CEN/TC 251, PT1-013, November 1995.
11. C. Chronaki, X. Zabulis, N. Stathiakis, E. Leisch, D. Katehakis, M. Tsiknakis, and S.C. Orphanoudakis. A Web-Based Environment for Medical Collaboration in a Regional Healthcare Network. In Proceedings of CAR'97, Berlin, Germany, June 25-28, 1997.
12. Committee on Information and Communications (CIC). High Performance Computing and Communications: Foundation for America's Information Future. National Science and Technology Council, Washington, D.C, 1995.
13. P. Constantopoulos and M. Doerr. The Semantic Indexing System. Technical Report, Institute of Computer Science, FORTH, Heraklion, Crete, Greece, 1995.
14. T.J. Mowbray and R. Zahavi. The Essential CORBA: Systems Integration Using Distributed Objects. OMG, 1995.
15. M. deBellis and S.B. Tantry. An Object-oriented Healthcare Information Infrastructure. In Proceedings of OOPSLA, 1995.
16. D.M. Dilts and W. Wu. Using Knowledge-Based Technology to Integrate CIM Databases. IEEE Expert, Vol. 3, No. 2, June 1991, pp. 237-245.
17. Digital Imaging & Communication in Medicine (DICOM), Version 3.0. American College of Radiologists, National Electrical Manufacturers Association, 1993.
18. D. Gunning. Intelligent Integration of Information. Presentation Slides, ARPA/ISO 1995 Annual Program Overview.
19. D.F. Ferguson, C. Nikolaou, J. Sairamesh, and Y. Yemini. Economic Models for Allocating Resources in Computer Systems. In: Market based Control of Distributed Systems, Ed. Scott Clearwater, World Scientific Press, 1995.
20. D. Forslund, R. Phillips, D. Kilman, and J. Cook. Experiences with a Distributed Virtual Patient Record System. In Proceedings of the AMIA'96 Fall Symposium.
21. D. Forsland. The Role of CORBA in enabling Telemedicine. In Proceedings of the Global Forum III in Telemedicine, Vienna, Virginia, March 1997.
22. D. Floresu, L. Raschid, and P. Valduriez. A Methodology for Query Reformulation in CIS using Semantic Knowledge. Inter. Journal of Cooperative Information Systems, V. 5, No. 4, (1996), 431-467, pp. 431-467.
23. J. Frew, M. Freeston, R.B. Jemp, J. Simpson, T. Smith, A. Wells, and Q. Zhang. The Alexandria Digital Library Testbed. D-Lib Magazine, July/August 1996.
24. D. Garlan. Research Directions in Software Architecture. ACM Computing Surveys, 27(2):258-261, June 1995.
25. D. Georgakopoulos et.al. An Overview of Workflow Management: From Process Modeling to Workflow Automation Infrastructure. Distributed and Parallel Databases, vol. 3, no 2, April 1995, pp. 119-153.
26. J. Hammer. An Intelligent Interface in an Associate System. In: Human/technology interaction in complex systems (Vol. 7), Greenwich CT: JAI Press.
27. S. Heiler. Semantic interoperability. ACM Computing Surveys, 27(2):271-273, June 1995.
28. HL7: An Application Protocol for Electronic Data Exchange in Healthcare Environments, Version 2.1. Health Level Seven Inc., Ann Arbor, Michigan, 1991.
29. Committee on Applications and Technology. Healthcare and the NII: Draft for public comment. WWW Document, 1994.

30. C. Hsu, M. Bouziane, W. Cheung, L. Rattner, and L. Yee. Metadatabase Modeling for Enterprise Information Integration. *Journal of Systems Integration*, Vol. 2, No. 1, 1992, pp. 5-37.
31. C. Hsu and L. Rattner. Metadatabase Solutions for Enterprise Information Integration Problems. *Data Base*, Vol. 24, No. 1, Winter 1993, pp. 23 - 35.
32. H.K. Huang and S.T.C. Wong. The Medical Informatics Infrastructure within an Integrated Hospital Environment. In *Proceedings of EuroPACS 96*, Heraklion, Crete, Oct. 3-5, 1996.
33. B.L. Humphreys, D.A.B. Lindberg, and W.T. Hole. Assessing and enhancing the value of UMLS Knowledge Sources. In *Proc. of the 15th Symp. on Computer Applications in Medical Care (SCAMP) 1991*, pp. 78-82.
34. V. Jagannathan, Y. Reddy, K. Srinivas, R. Karinithi, R. Shank, S. Reddy, G. Almasi, T. Davis, S. Qiu, S. Friedman, B. Merkin, and M. Kilkenny. An Overview of the CERC ARTEMIS project. Technical Report CERC-TR-RN-95-002, Concurrent Engineering Research Center, West Virginia University, April 1995.
35. N.R. Jennings, P. Faratin, M.J. Johnson, and T.J. Norman. Agent-based Business Process Management. *International Journal of Cooperative Information Systems*, 4(2& 3):105-130, 1996.
36. D. Katehakis, M. Tsiknakis, A. Armaganidis, and S.C. Orphanoudakis. Functional and Control Integration of an ICU, LIS and PACS Information System. In *Proceedings of MIE'97*, Porto Carras, Greece, May 25-29, 1997.
37. J.A. Kunze and B.N. Warling. Recent Developments in GALEN II: Evolution of a digital library for the Health Sciences. *D-Lib Magazine*, March 1996.
38. C. Lagoze. The Warwick Framework: A Container Architecture for Diverse Sets of Metadata. *D-Lib Magazine*, July/August 1996.
39. D.E. O'Leary. The Internet, Intranets, and the AI renaissance. *IEEE Computer*, pages 71-78, January 1997.
40. E. Leisch, M. Tsiknakis, and S.C. Orphanoudakis. An Architectural Framework for the Integration of Geographically Distributed Heterogeneous Autonomous Medical Information Systems. In *Proceedings of EuroPACS'96*, Heraklion, Crete, Greece, pp. 73-77, October 3-5, 1996.
41. E. Leisch, M. Tsiknakis, and S.C. Orphanoudakis. Pre-Hospital Health Emergency Management as an Integrated Service of the Regional Health Telematics Network of Crete. In *Proceedings of MIE'97*, Porto Carras, Greece, May 25-29, 1997.
42. F. Manola. Interoperability Issues in Large-scale Distributed Systems. *ACM Computing Surveys*, 27(2):268-270, June 1995.
43. D.C. Martin, C.S. Ang, M.E. Solomon, and M.D. Doyle. Libraries in the Information Age. Technical report, UCSF Library and Center for Knowledge Management, 1993.
44. D.G. Messerschmitt. What the NII could be: A User Perspective. White paper, Department of Electrical Engineering and Computer Sciences, 1995.
45. D.G. Messerschmitt. The Convergence of Telecommunications and Computing: What are the implications today? *Proceedings of the IEEE*, 84(8), August 1996.
46. N. Natarajan and G.M. Slawsky. A Framework Architecture for Multimedia Information Networks. *IEEE Communications Magazine*, February 1992, pp. 97 - 104.
47. O. Nierstrasz and D. Meijler. Research Directions in Software Composition. *ACM Computing Surveys*, 27(2):262-264, June 1995.
48. A.T. McCray and A. Razi. The UMLS Knowledge Source Server. *MEDINFO 1995*, pp 1447-1456, 1995.
49. S.C. Orphanoudakis, C.E. Chronaki, and D. Vamvaka. *I²Cnet*: Content-Based Similarity Search in Geographically Distributed Repositories of Medical Images. *J. Computerized Medical Imaging and Graphics*, vol. 20(4), pp. 193-207, 1996.
50. S.C. Orphanoudakis, E. Kaldoudi, and M. Tsiknakis. Technological advances in teleradiology. *European Journal of Radiology*, vol. 22, pp. 205-217, 1996.
51. A. Petru. Patient Meta Record - Global Data Model. Internal Report. CMI-HTA, FORTH, April 1997.
52. A. Paepcke, S. Cousins, H. Garcia-Molina, S. Hassan, S. Ketchpel, M. Roscheisen, and T. Winograd. Using distributed objects for digital library interoperability. *IEEE Computer Magazine*, 29(5):61-68, May 1996.
53. A. Paepcke. Digital libraries: Searching is not enough; what we learned on-site. *D-Lib Magazine*, May 1996.
54. Y. Papakonstantinou, H. Garcia-Molina and J. Widom. Object Exchange across Heterogeneous Information Sources. In *Proceedings of IEEE International Conference on Data Engineering*, 1995.
55. D.V. Pargaonker, V. Jagannathan, Y. Al-Salqan, and T. Davis. Intelligent Agents, HL7 and Patient Record Databases. Technical Report CERC-TR-RN-95-003, Concurrent Engineering Research Center, West Virginia University, 1995.
56. R.L. Phillips. A Network-based Distributed, Media-rich Computing and Information Environment. In *Proceedings of Digital Media and Electronic Publishing*, Leeds, UK, 6-8 December, 1994, December 1994.
57. X. Qian, G. Wiederhold, M. Bilello, A. Chavez, and V. Sarathy. Trusted Interoperation of Healthcare Information. In *NSF Challenge workshop*, March 1996.
58. R. Rao. Rich Interaction in Digital Libraries. *Communications of the ACM*, April 1995, Vol. 38, No. 4, pp. 29-39.

59. A. Sanchez. User Agents in the Interface to Digital Libraries. In Proceedings of DL'94.
 60. J. Sairamesh, C. Nikolaou, D. Ferguson and Y. Yemini. An Economic Framework for Pricing and Charging in Digital Libraries. *Dlib Magazine*, February 1996.
 61. E. Sciore, M. Siegel, and A. Rosenthal. Using Semantic Values to Facilitate Interoperability Among Heterogeneous Information Systems. *ACM Transactions on Database Systems*, Vol. 19, No. 2, June 1994, pp. 254-290.
 62. S. Shen, R. Makkamala, A. Wadda, C. Zang, H. Abdel-Wahab, K. Maly, A. Liu, and M. Yuan. An Interoperable Architecture for Digital Information Repositories. In Proceedings of Digital Libraries'94, 1994.
 63. K. Sycara and D. Zeng. Coordination of Multiple Intelligent Software Agents. *International Journal of Cooperative Information Systems*, 5(2& 3):181-211, 1996.
 64. Telemed Team. Telemed Virtual Patient Record Systems. Web document, Los Alamos National Laboratory, 1995.
 65. J.T.C. Teng and W.J. Keffinger. Business Process Redesign and Information Architecture: Exploring the Relationships. *Data Base*, Vol. 26, No. 1, February 1995, pp. 30-42.
 66. R.M. Tong and D.H. Holtzman. Knowledge-based access to Heterogeneous Information Sources. In Proceedings of DL'94.
 67. M. Tsiknakis, D. Katehakis, and S.C. Orphanoudakis. Intelligent Image Management in a Distributed PACS and Telemedicine Environment. *IEEE Communications Magazine*, vol. 34, no. 7, July 1996, pp. 36-45.
 68. M. Twindale and D.M. Nichols. Interfaces to Support Collaboration in Information Retrieval. In Proceedings of the BCS IR & HCI Workshop (Eds.) Johnson, C. & Dunlop, M. 1996, Glasgow, UK, 25-28.
 69. Visible Human Project. National Library of Medicine. <http://www.nlm.nih.gov>.
 70. S. Weiber and C. Lagoze. An Element Set to Support Resource Discovery. Technical report, Computer Science Department, Cornell University, January 1997.
 71. G. Wiederhold. Effective information transfer for healthcare, quality versus quantity. In NII 2000 Workshop Notes, Computer Science and Telecommunications Board, National Academy of Sciences, Washington DC, May 1995.
 72. G. Wiederhold. Mediation in Information Systems. *ACM Computing Surveys*, 27 (2):265-267, June 1995.
 73. G. Wiederhold, M. Bilello, V. Sarathy, and X. Qian. A Security Mediator for Healthcare Information. In Proceedings of the American Medical Informatics Association Annual Fall Symposium, pages 120-124, 1996.
 74. S.T.C. Wong and H.K. Huang. A Hospital Integrated Framework for Multimodality Image Base Management. *IEEE Trans. On Systems, Man, and Cybernetics*, Vol. 26, No. 4, pp 455 - 469, 1996.
 75. Cross Industry Working Team. An Architectural Framework for the National Information Infrastructure. White paper, September 1994. Available from Corporation for National Research Initiatives, Reston, Virginia (CNRI).
 76. Cross Industry Working Team. Class Profiles for the Current and Emerging NII. White paper, February 1997. Available from Corporation for National Research Initiatives, Reston, Virginia (CNRI).
 77. M. Zikos, G. Stephanidis and S.C. Orphanoudakis. CoMed: Cooperation in Medicine. In Proceedings of EuroPACS '96, Heraklion, Greece, 1996, pp.88-92.
 78. M. Zikos, E. Kaldoudi, and S.C. Orphanoudakis. Image Processing within an Integrated Teleradiology Services Network. In Proceedings of CAR'97, Berlin, Germany, June 25-28, 1997.
-
- Manolis Tsiknakis** received a B. Eng. degree in electronic engineering, a MSc. in micro-processor engineering, and a Ph.D. in control systems engineering from the University of Bradford, U.K. Dr Tsiknakis was a member of the research staff at the department of Control Systems Engineering at the University of Bradford from 1985 till 1990. In 1992 he joined the Institute of Computer Science of the Foundation for Research and Technology - Hellas (ICS-FORTH), where he is currently coordinating the activities of the Center of Medical Informatics and Health telematics Applications. His current research interests are in the areas of: multimedia systems with emphasis on medical information systems, interoperability in heterogeneous systems, information integration, Telematics Engineering, and Signal Processing and Analysis.
- Catherine Chronaki** holds a B.S. in Computer Engineering from University of Patras, Greece, and a MSc in Computer Science from Rochester Institute of Technology and University of Rochester. In 1992, she joined the Center for Medical Informatics and Health Telematic Application, in the Institute of Computer Science, FORTH, in Crete. Her research interests include but are not limited to medical information systems, Internet applications, integration architectures, and digital medical libraries.
- Sarantos Kapidakis** holds a Diploma in Electrical Engineering from NTUA and a MSc. and a Ph.D. degree in Computer Science from Princeton University. In 1993, Dr Kapidakis joined the Computer Science Dept of the University of Crete and the Institute of Computer

Science of the Foundation for Research and Technology - Hellas, where he is currently working on the Parallel and Distributed Systems Group, coordinating the Digital Libraries Activities. His current research interests include algorithms, programming languages, operating systems, digital libraries, electronic commerce and parallel and distributed systems in general.

Christos Nikolaou received his Diploma in Electrical Engineering from the National Technical University of Athens in 1977, the M.S. degree in Applied Mathematics from Harvard University, Cambridge, in 1979, and the Ph.D. degree in Applied Mathematics from Harvard University in 1982. He is the Vice-Rector for Financial Planning and Development and an associate professor of computer science at the University of Crete, Greece. At the Institute of Computer Science of FORTH, he is the head of the Parallel and Distributed Systems Group (Pleiades). He was a Research Staff Member at the IBM T.J. Watson Research Center from 1983 to 1992. He has many years of academic and research experience in the fields of distributed systems and databases, and is a leading participant of the ERCIM Digital Library Initiative and the ESPRIT DELOS Digital Library Working Group. He has served on various committees and working groups of the European Commission and has been active in European R&D programs. He currently serves as the Chairman of the Executive Committee of the European Research Consortium for Informatics and Mathematics (ERCIM).

Stelios C. Orphanoudakis received the B.A. degree in engineering sciences from the Dartmouth College, Hanover, New Hampshire, in 1971, the M.S. degree in electrical engineering from Massachusetts Institute of Technology, Cambridge, in 1973, and the Ph.D. degree in electrical engineering from Dartmouth College in 1976. He is director of the Institute of Computer Science, Foundation for Research and Technology-Hellas, and professor of computer science, University of Crete, Greece. He held a faculty appointment in the Departments of Diagnostic Radiology and Electrical Engineering at Yale University from 1975 until 1991. Professor Orphanoudakis has many years of academic and research experience in the fields of computer vision and robotics, intelligent image management and retrieval by content, and medical imaging. He has served on various committees and working groups of the European Commission and has been active in European R&D programs. He currently

serves on the board of Directors and is Vice President of the European Research Consortium for Informatics and Mathematics (ERCIM). He is also a member of the National Telecommunications Commission and the National Advisory Research Council of Greece.