Delivering a Lifelong Integrated Electronic Health Record Based on a Service Oriented Architecture

Dimitrios G. Katehakis, *Member, IEEE*, Stelios G. Sfakianakis, Georgios Kavlentakis, Dimitrios N. Anthoulakis, and Manolis Tsiknakis, *Member, IEEE*

***Abstract*—Efficient access to a citizen's Integrated Electronic Health Record (I-EHR) is considered to be the cornerstone for the support of continuity of care, the reduction of avoidable mistakes, and the provision of tools and methods to support evidence-based medicine. For the past several years, a number of applications and services (including a lifelong I-EHR) have been installed, and enterprise and regional infrastructure has been developed, in HYGEIAnet, the Regional Health Information Network (RHIN) of the island of Crete, Greece. Through this paper, the technological effort toward the delivery of a lifelong I-EHR by means of World Wide Web Consortium (W3C) technologies, on top of a service-oriented architecture that reuses already existing middleware components is presented and critical issues are discussed. Certain design and development decisions are exposed and explained, laying this way the ground for coordinated, dynamic navigation to personalized healthcare delivery.**

***Index Terms*—Integrated Electronic Health Record, Regional Health Information Network, service-oriented architecture, Web Services.**

I. Introduction

THE catalyst for change in the health sector, based on the use of Information and Communication Technologies (ICTs), is the need for improved quality of healthcare services and the containment of related costs. As citizens become better educated and informed, they increasingly have higher expectations regarding access to care and the quality of healthcare services. In the future, healthcare professionals will continue to deliver care, but will also be increasingly required to share their knowledge and expertise with other colleagues, while citizens will demand that they actively participate in medical decisions. Allowing access to services for patients will require significant organizational commitment at the early stages in order to improve the quality and usability of information. In this context, services and information must be accessible without (visible) organizational boundaries in order to support seamless and personalized information delivery.

The Integrated Electronic Health Record (I-EHR) of an individual consists of a collection of lifetime health data in electronic format, generated during relevant interactions with the healthcare system [1]. In addition to providing support for continuity of care, the I-EHR is expected to be a valuable tool in clinical research, medical decision-making, epidemiology, evidence-based medicine, and in formulating public health policy.

Anticipated benefits of an I-EHR include:

1) Availability and accessibility of vital health information 24 h a day, seven days a week, regardless of where the person requiring care happens to be.
2) More effective and efficient treatment since healthcare practitioners will be better positioned to spend more quality time with their patients.
3) Reduction of the number of redundant procedures, and therefore, less health risks for the patient and greater cost savings.
4) Empowerment of individuals to exercise greater control over their own health, by giving them access to their own personal health records, and by enabling them to make informed choices about options available to them.
5) Improved quality of care, as a result of the formulation of relevant healthcare policies, by means of collectively anonymized information contained within individual Electronic Health Records (EHRs).

Today, a large number of healthcare organizations offering similar services exist within the boundaries of any Regional Health Information Network (RHIN). Because of their competing interests, they retain their independence, and collaboration with each other is usually determined by their interests. Top priority of any RHIN is to make data and information securely available in the inter-enterprise environment where needed, when needed, and in the format needed. Within such an environment, the need for a single I-EHR for every citizen becomes the cornerstone for supporting continuity of care and providing content to evolving health telematics and e-health services.

Today, Web Services (WS) provide an open and standardized way to achieve interoperation between different software applications, running on a variety of platforms and/or frameworks. Therefore, they constitute an important technological tool toward the incremental delivery of advanced inter-enterprise services.

This paper begins with a brief overview of current activities related to EHR work conducted internationally (Section II). Section III focuses on the architecture required to be in place in order to support efficient service development across cooperating healthcare enterprises that are part of a RHIN. Section IV focuses on the building blocks an I-EHR consists of. The HYGEIAnet I-EHR, developed on the RHIN of Crete, Greece, is presented in Section V, where significant effort has been paid

Manuscript received August 22, 2005; revised April 7, 2006. This work was supported in part by several national and EU co-funded RTD projects.

D. G. Katehakis, S. G. Sfakianakis, G. Kavlentakis, and M. Tsiknakis are with the Institute of Computer Science, Foundation for Research and Technology-Hellas, GR-71110 Heraklion, Crete, Greece (e-mail: katehaki@ics.forth.gr; ssfak@ics.forth.gr; gkab@ics.forth.gr; tsiknaki@ics.forth.gr).

D. N. Anthoulakis is with Usableweb Ltd., 71601 Heraklion, Crete, Greece (e-mail: d.anthoulakis@usableweb.gr).

Digital Object Identifier 10.1109/TITB.2006.889711

on utilizing the already existing (legacy) software components of the health information infrastructure (HII), based on the Common Object Request Broker Architecture (CORBA) technology, in order to deliver the I-EHR as a WS. Finally, Section VI concludes by presenting prospective issues and future trends.

II. Background

So compelling are the arguments for an I-EHR service that a number of countries around the world—including the United States, Canada, several European Union (EU) countries, and Australia—are striving to develop workable models along with their vision for their national HII.

At the end of 2003, the Medicare Prescription Drug Improvement and Modernization Act (MMA) was signed in the United States including, among other new initiatives, important provisions for Health Information Technology. MMA requires the Centers for Medicare and Medicaid Services to develop standards for electronic prescribing, which is expected to be a first step toward the widespread use of EHRs. In addition, the MMA requires the establishment of a Commission on Systemic Interoperability to provide a road map for interoperability standards in order for the majority of the U.S. citizens to have interoperable electronic health records within ten years.

At the same time, the American Society for Testing and Materials has established a standard for the Continuity of Care Record (CCR) addressing more directly the issue of patient data summaries used for transfers, referrals, and discharges. The CCR includes basic, minimum data such as diagnoses, procedures, medications, and care plans, and is seen as an intermediate, short-term solution to an interoperable EHR system, and can be implemented as a Health Level Seven (HL7) Version 3 (V3) Clinical Document Architecture (CDA) eXtensible Markup Language (XML) document.

In Canada, the Canadian Infoway, an independent, nonprofit corporation, initiated as a result of the Canadian federal government's announcement in September 2000 to accelerate the development and adoption of modern systems of Information Technology (IT) in healthcare, aims to foster the development of secure and interoperable EHR systems across Canada. Its objectives are to develop mechanisms to enable consumers to access health information they can use to facilitate the work of healthcare providers through technology, and to create a unified network of EHRs across the continuum of care.

In Europe, the U.K. National Health Service (NHS) information strategy for health represents an evolutionary approach to the development of an I-EHR service. Today, the "National Programme for IT in the NHS" is trying to hold electronically every patient's medical and care records, and eventually make it available securely online, starting from England. For doing so, it uses contracted service providers for the implementation and operation of the new technology for the NHS care record service and related information services.

In Denmark, shared information is considered as the foundation for seamless care and patient involvement, and therefore, emphasis is placed on the development of common terminology and classifications. Subsequently, a major initiative directly linked to the 2003–2007 national strategy is related to the coordinated development, test, and implementation of EHRs based on the elements clinical processes consist of.

In Germany, the central associations of self-administration committed themselves in 2002 "to develop a new infrastructure for telematics on the basis of a general framework architecture, to improve and/or introduce the electronic communication (electronic prescription, electronic discharge letter by the physician) and to introduce the former health insurance card as an electronic health card in the future" [2]. The target is to have 80 millions electronic health cards distributed in 2006 that will be capable of giving access to medical data, and therefore, leadin to the Electronic Patient Record.

Australia is taking an incremental approach, focused on starting systems in conjunction with provincial and state governments. HealthConnect, the Australian EHR project, backed by the country's Department of Health and Ageing requires the use of standardized electronic clinical messages to support enhanced clinical communications among healthcare providers, built around patient information summaries, not full clinical records.

Regarding international standardization efforts, the International Organization for Standardization Technical Committee 215 (ISO/TC215), in its technical report 20514 (Electronic Health Record Definition, Scope, and Context) [3] and its technical specification 18308 (Requirements for an EHR Architecture) [4], provides the definition of EHR, and delivers a consolidated set of EHR requirements for using, sharing, and exchanging EHRs, independent of technology and current organization structures. The European Committee for Standardization Technical Committee 251 (CEN/TC 251) has started revising European Prestandard (ENV) 13606 [5]–[8] to provide a rigorous and durable information architecture for communicating EHR to support the interoperability of systems and components interacting with EHR services, by having adopted the OpenEHR methodology and by using ISO 18308 as an EHR architecture standard. The revised CEN European Standard (EN) 13606 will also include compliance with HL7 CDA Release 2. At the same time, HL7 has approved the Electronic Health Record System (EHR-S) Functional Model to move forward as a Draft Standard for Trial Use (DSTU) intending to provide a summary understanding of functions that may be present in an EHR-S, from a user perspective, in order to enable consistent expression of system functionality.

III. RHIN Architecture

Any RHIN can be considered as an extended virtual healthcare enterprise, where several e-health services, different in nature, exist, and are available among trading/collaborating partners/enterprises.

The approach that is about to be described does not focus on the standardization of the content of the EHR, but it rather tries to provide access to already existing EHR systems. As such, the federated approach was selected, based on a set of already

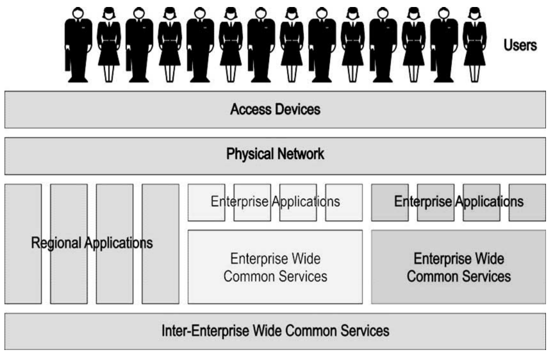

Fig. 1. Architecture of a RHIN.

available regional HII components, as described in the EU co-funded RTD Project "Professionals and Citizens Network for Integrated Care" (PICNIC) [9] that dealt with inter-enterprise integration issues, and has developed a model for the future RHINs. According to PICNIC, the architecture of a regional network of healthcare services can be presented as comprising a number of layers (see Fig. 1).

Basic characteristics of the PICNIC architecture include the following:

1) users (healthcare professionals and citizens) that have access to applications and services (either enterprise or regional) through a secure physical network by means of alternative access devices;
2) enterprise applications that are supported by enterprise-wide middleware services;
3) regional applications and services that are supported by the corresponding common inter-enterprise wide services (regional health infrastructure).

This type of multi-tier approach depends heavily on the existence of both generic and healthcare specific middleware services/components, and imposes a level of common design that varies according to the actual composition of the platform. Healthcare-related components are needed for the proper identification of the subjects of care, the exchange of I-EHR indexing and health data (utilizing appropriate health-oriented protocols like HL7), health resource(s) location(s), facilitation of collaboration between healthcare professionals and patients/experts, authorization for accessing healthcare-related resources, and the management of common medical terminology. Generic components are required to support low-level, essential, platform-dependent functionalities like concurrency control, notification/event handling, licensing, security (authentication, encryption, auditing, etc.), timing, transaction management, etc.

In practice, the architecture of a RHIN is likely to be much more complex. This is because healthcare providers in a region have applications running, supporting their internal operations, together with intra-enterprise-wide integration platforms of messaging and common IT services. In addition, the functionality provided by the RHIN depend on the IT services that organizations desire to provide to each other and what they *are willing to give up*, in order to achieve tighter coupling. Therefore, the common IT services infrastructure depends on the already existing HIIs of the healthcare organizations that make up the RHIN of providers and related organizations.

The IEEE Std. 1471-2000 [10] provides the necessary guidelines for the description of a RHIN architecture since it addresses the concerns of the identified stakeholder groups, consisting of the citizens, the regional health planners, the service providers, the users, the designers, the developers, the implementers, the maintainers, and the involved vendors. According to the standard, there exist multiple viewpoints for the description of the architecture of an integrated system.

When the system under consideration is a RHIN, besides its Conceptual Viewpoint, there is a need for addressing the following viewpoints, as well:

1) the *Enterprise Viewpoint*, which focuses on the purpose, goal, and policies governing the integrated RHIN applications and services;
2) the *Information Viewpoint*, which focuses on the analytical presentation of how semantic homogeneity of information is achieved, and what concepts and vocabularies are necessary for the implementation of the RHIN;
3) the *Computational Viewpoint*, which focuses on the description of the subsystems the overall system consists of, their functionalities, and the precise definition of the subsystem interfaces (encapsulation of capabilities, separation of functionalities);
4) the *Engineering Viewpoint*, which focuses on the interaction between the components of the overall system;
5) the *Technology Viewpoint*, which focuses on the description of the physical implementation of the system;
6) the *Security Viewpoint*, which focuses on how the system handles issues related to security.

IV. RHIN Infrastructure for the I-EHR

The realization of an I-EHR depends on the ability to provide integrated access to various parts of one's EHR and not necessarily to physically integrate them. Effectively, what is required is to keep the various healthcare enterprises self-contained and autonomous, but still put in place the mechanisms and the infrastructure that will support continuity of care, based on a federated model that can be applied on demand. Any such approach should be capable of providing uniform ways for accessing authentic, physician-generated, EHR information that is physically located in different clinical information systems within the boundaries of the healthcare enterprises the RHIN consists of. Furthermore, it must be able to provide fast and authorized on-line access to longitudinal views of an individual's personal health record to allow for the timely delivery of healthcare.

The I-EHR should not be confused with autonomous clinical information systems, message-based communication of EHR data, centralized clinical data repositories, portable segments of one's EHR, or monolithic information systems that have embedded in their structure mechanisms for accessing directly host systems.

The main challenges that need to be effectively addressed, in an effort to create the HII required to support a federated I-EHR, include the following:

1) the definition/adoption of a federated/global reference schema for the I-EHR that is capable of supporting and providing effective solutions to immediate needs within the RHIN, without imposing significant constraints with the incorporation of new systems in the federation;
2) the use of standardized interfaces for accessing EHR information, either directly by the end user or through a middleware set of components managing the required minimum data sets, as well as indexing;
3) the population of the individual components of the software infrastructure by the corresponding, domain-specific, workable, and semantically consistent concept models to be used for communicating structures containing values of data;
4) the adoption of common terminology and a common naming convention for the use of globally unique concept codes formed by combining accepted coding scheme identities and the local concept code within that coding scheme;
5) the use of available technology to support standardized means of communication among different software applications and services;
6) the implementation of the required mechanisms to enable information readiness, consistency to guarantee the required quality of service;
7) the implementation of an adequate security system, which will deliver an I-EHR service that its users can trust, with consent management being part of the overall security policy.

A. I-EHR Content

When trying to develop an EHR system highly adaptable to future requirements, a clear separation between knowledge and information is required. Such an EHR system must be based on a reference model that will be in a position to endure. At the same time, ephemeral needs must be covered through a clearly specified domain model, capable of working by means of a controlled vocabulary, while certain needs for external communications must be provided through well-defined interfaces (see Fig. 2).

In order to provide access to already available information, outside the strict boundaries defined by each individual EHR system, the standardized interface that has to be defined must be based on a common/federated domain model, a standardized EHR reference model, and a corresponding controlled common/federated vocabulary. Subsequently, part (or the whole) of the schema of each information system participating in the federation must be in a position to be mapped to that particular common, normalized schema (the federated schema).

Since the I-EHR has not been defined to be the EHR itself but rather a service through which integrated access to EHR information is provided, it must be capable of resolving difficult indexing and location issues. Access to potentially any kind of EHR information must be supported, including text and numeric values, structured (e.g., HL7 CDA) and unstructured documents, as well as multimedia information like waveforms, sound, and image files.

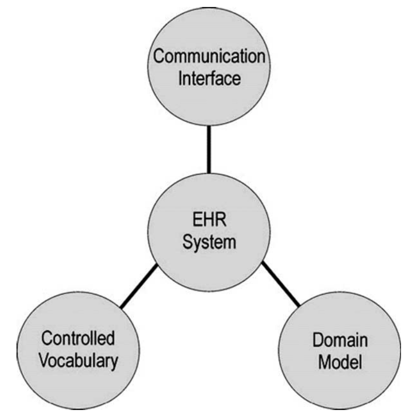

Fig. 2. A typical interoperable EHR system must be capable of supporting an ephemeral domain model to enable future domain model extensions, support interoperability without any dependence on external terminology through controlled vocabularies, and a communication interface to the outside world.

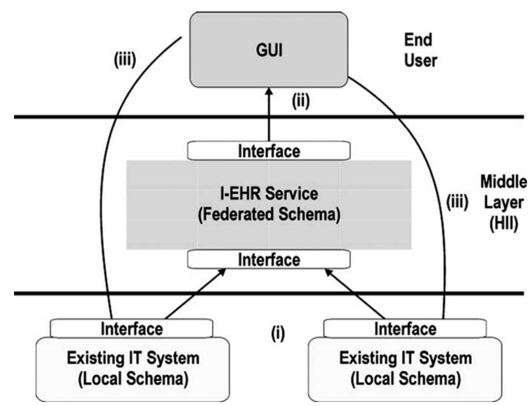

Fig. 3. In order for the end user to access EHRs, (i) existing IT systems propagate up-to-date information to the middle layer of the infrastructure, (ii) the end user retrieves indices, selects information of interest, and (iii) directly accesses clinically significant information.

B. Conceptual Approach—Components of the Infrastructure

The I-EHR must provide fast, secure, and authorized access to distributed patient record information from multiple, disparate sources, and an environment for integrated, round-the-clock access to clinically significant information that is kept at the place where it is produced and acquired by the most appropriate clinical information system. As seen in Fig. 3, elements that are necessary at a RHIN setting for the delivery of an I-EHR involve:

1) information propagation from feeder systems to the middle layer of the HII;
2) a component residing at the middle layer managing the required minimum data sets, as well as indexing;
3) a Graphical User Interface (GUI) to make available the I-EHR service to the end users.

Main actors identified to be related with the I-EHR service include:

1) the *End User* that represents the ultimate user of the system and can be either a *Citizen* or a *Healthcare Professional* and expects the service to be able to offer role-based, secure access to reliable, patient information 24 hours a day;

TABLE I
USE CASES FOR THE I-EHR IT SERVICE

Use Case	Short Description
Add New IT System	It adds a new clinical information system to the federation of IT systems. The *Maintainer* initiates this use case.
Access Control Rule	It defines, sets, and updates the security roles and accesses policies of the users. The *Maintainer* or the *End User* initiates this use case.
Keep Information Up-to-date	Accesses the internal data structures so that they are kept up-to-date. It is initiated by the *Maintainer* or by an *Existing IT System*. It uses three sub-use cases: "Provide Secure Information Communication", "Keep Auditing Trails", and "Semantic Mapping".
Access Rights Authentication	It relates to authentication and the use of passwords and user groups to validate users and allow them to access clinical information.
Access Clinical Information	It is initiated by the *End User* to access clinical information. It uses three sub-use cases: "Provide Secure Information Communication", "Keep Auditing Trails", and "Semantic Mapping".
Provide Secure Information Communication	It provides security through all communication paths. It is used by the "Access Clinical Information" and the "Keep Information Up-to-date" use cases.
Keep Auditing Trails	It records every interaction between the various entities of the system. It is used by the "Access Clinical Information" use case and the "Keep Information Up-to-date" use case.
Semantic Mapping	Performs the translation between languages and coding schemes. It is used by the "Access Clinical Information" and the "Keep Information Up-to-date" use cases.
Add New IT System	It adds a new system to the federation. The *Maintainer* initiates this use case.

2) the *Maintainer* that represents a person responsible for the maintenance, expansion, and enhancement of the system with new features;
3) the *Existing IT System* that represents a system that is the source of primary clinical information.

A set of I-EHR features that can be translated into use cases is listed in Table I, and the corresponding Unified Modeling Language (UML) diagram is depicted in Fig. 4.

The I-EHR environment, as has been described, provides a decentralized view of the patient medical record, by dynamically composing information that resides in a variety of heterogeneous, self-consistent, clinical information systems that have been optimized with respect to the requirements of different medical specialties and levels of care. The initial sets of essential HII services that have been identified as required include:

1) patient identification services for the unique identification of patients, based on a commonly agreed set of traits, and for the correlation of identities across different identification domains;
2) I-EHR indexing services for the provision of necessary means for the location of clinically significant information throughout the participating enterprises' clinical information systems;
3) clinical observations access services for direct access to the sources of clinical information, where the complete, original (physician-generated, or acquired from a medical device) EHR data is kept;
4) I-EHR update-brokers for the provision of prompt and consistent propagation of indexing information to the I-EHR indexing service. Indexing information can be either extracted by an I-EHR update-broker through a clinical observation access service (pull model), or forwarded by locally installed clinical observation access services to the I-EHR update-broker (push model);
5) health resource services for the unique identification of related resources such as organizations, devices, and/or software and the means for accessing them;
6) terminology services for the association of coding schemes in use and to enable the transformation of information from one form or representation to another;
7) security services (e.g., for auditing, encryption, authentication, etc.) under a common public key infrastructure framework to counter all kinds of security threats;
8) user profile services for maintaining personalized settings and preferences.

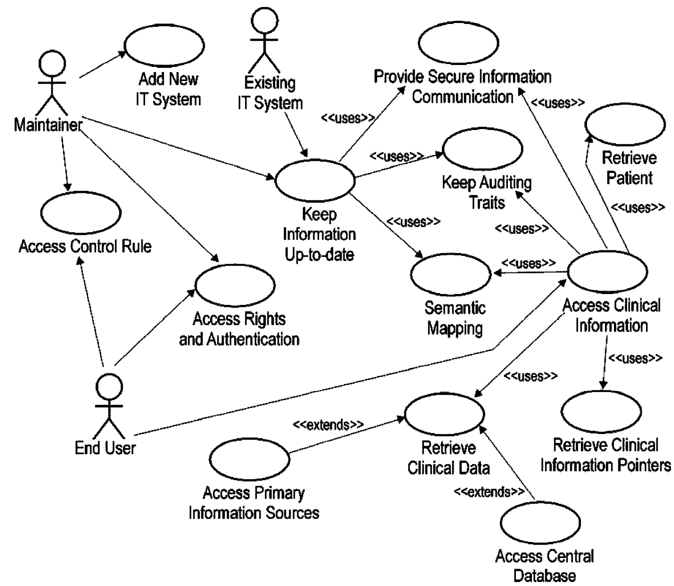

Fig. 4. Main use case diagram for the I-EHR service.

The end user, through the patient identification service, identifies the patient under consideration and accesses indexing information through the I-EHR indexing service. Once the location becomes known, the end user uses the corresponding clinical observation access service to access clinical data.

C. Implementation Approach—Why WS?

WS is a major, service-oriented, connection technology, which is specification based and mostly open. In addition to its open source development potential in a technology neutral environment, major vendors are embracing World Wide Web Consortium (W3C) and the Internet Engineering Task Force (IETF) efforts.

Significant advantages of using WS on top of any existing middleware solution is location transparency, language and

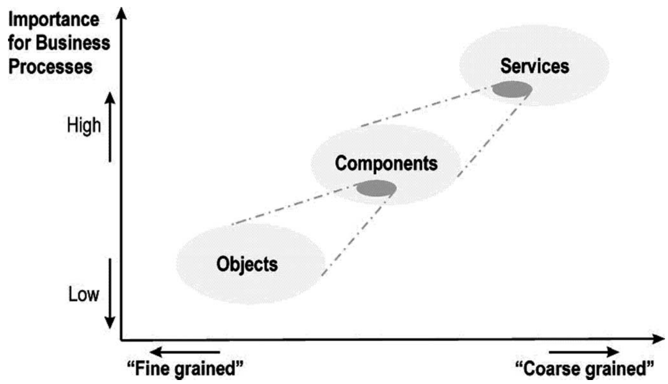

Fig. 5. Evolution of middleware technologies.

platform independence, together with their embracement by big vendors and the acceptance they enjoy between the users. The popularity of WS, and the excitement that they evoke to enterprise users, ease the adoption of relevant applications and frameworks. On the other hand, legacy technologies now rest in assurance that they are technologies that simply work within an enterprise environment despite the various shortcomings. The more important middleware players Java Two Enterprise Edition, CORBA, and (Distributed) Component Object Model, are strong and each one enjoys a big share of the installed application base.

WS, with their XML roots, open the door for the next generation, loosely coupled, coarse-grained, document-oriented architectures (see Fig. 5). The term "loosely coupled" is used to characterize services where the assumptions made by the communicating parties about each other's implementation, internal structure, and the communication itself are kept to a minimum. The expected result is systems that are resilient to change and flexible enough so that they can evolve independently. Asynchronous interaction models are generally better suited in this regard, since they free the sender and receiver parties from a commitment to a predetermined dialog agreed in a contract, as is the case in synchronous Remote Procedure Call (RPC) like environments. The asynchronous communication paradigm naturally leads to event-driven exchange of messages and departs the procedural RPC model. Message queues and relevant technologies perfectly fit in these scenarios, where documents are exchanged using a well-defined and general interface (the send/receive message methods or equivalent) and the importance shifts from the definition of interface contracts to the design of the document schemas.

The adoption of WS assumes the satisfaction of various horizontal requirements in existing middleware platforms, like the following.

1) *Communication protocol interoperability:* Version 1.2 of the Simple Object Access Protocol (SOAP), which is the serialization format of the XML messages on the "wire," is a major effort to overcome some of the interoperability obstacles produced by the earlier versions of this protocol. Furthermore, specifications such as these produced by the WS Interoperability Organization (WS-I) help in formulating the minimal set of criteria so that two independent implementations could communicate with each other.
2) *Secure and auditable conversations in accordance with trust establishment policies:* Security should not be considered an afterthought but it ought to be built into the communication platform itself. WS were originally marketed as an easy way to do business across the Internet, based on their tunneling through the hypertext transfer protocol that usually bypasses corporate firewalls. The use of transport layer security is, most of the times, not enough to provide the desired levels of authentication, authorization, and trust. Nowadays, the exploitation of technologies like XML-Signature, XML-Encryption, and WS-Security is considered mandatory in order to achieve the necessary quality of protection for message integrity and confidentiality. Additional efforts such as WS-Trust, WS-Policy, and WS-SecureConversation are under way.
3) *Naming and discovery services:* Locating WS given certain functional, qualitative, or other criteria is another important aspect. The universal description, discovery, and integration registry, together with the electronic business XML (ebXML) registry constitute two of the most prevalent technologies.

More specialized features such as support for transactions, business processes, asynchronous event-based information communication, etc., are specified and start to become available both as tools and services. It is envisioned that the convergence of the XML document-oriented WS with the messaging-oriented middleware will be a turning point in the deployment of the inter-enterprise interaction scenarios over the Web.

D. Security Considerations

The focus in dealing with I-EHR security is related to the control of information, especially as it deals with private and confidential patient data. Security within an information domain is always established in accordance with the respective security policy. For communication between different information domains, a trusted end-to-end communication policy must be established.

ISO/TC 215, in its technical report 21089 (trusted end-to-end information flows) [11], offers a guide to trusted end-to-end information flow for EHRs and to the key trace points and audit events in the electronic entity/act record lifecycle (from point of record origination to each ultimate point of record access/use). The main concept is that it is the responsibility of the information domains to negotiate under what terms they are able and willing to exchange information. In general, these terms of collaboration are called access rights and can be managed on two levels:

1) *Authentication*, as the process of ensuring that the communicating party is the one claimed to be;
2) *Authorization*, as the process of ensuring that the communicating party is eligible to request for a specific action.

In addition to the access rights that govern the identity and eligibility of communicating parties, audit trails are needed to ensure accountability of actions of individual persons or entities, such as obtaining informed consent or breaching confidentiality. These records can be used to reconstruct, review, and examine transactions from inception to output of results. The records can also be used to track system usage, control authorized users, detect and identify intruders.

Currently, the most common technological tool to cover various security aspects is the Public Key Infrastructure (PKI). PKI is used to describe the processes, policies, and standards that govern the issuance, maintenance, and revocation of the certificates, public, and private keys that the encryption and signing operations require. PKI incorporates the necessary techniques to enable two entities that do not know each other to exchange information using an insecure network such as the Internet.

In order to guarantee the authenticity of a set of input data the same way a written signature verifies the authenticity of a paper document, PKI uses digital signatures.

Building on the digital signature technology, the digital signing of XML-based clinical documents is a special instance where the nature of the clinical workflow may require that each participant only signs that portion of the document for which he/she is responsible. Older standards for digital signatures do not provide the syntax for capturing this sort of high-granularity signature or mechanisms for expressing which portion a party wishes to sign.

V. The HYGEIAnet I-EHR

In the course of designing and implementing HYGEIAnet [12] in Crete, Greece, special efforts have been made to meet the requirements of the various user groups involved, and to use state-of-the-art technology and standards at every stage of development. Alternative patient, location, and problem-oriented views for the I-EHR have been considered in an attempt to provide transparent access and secure communication of information between medical specialty areas, as well as in a variety of situations from community to hospital care across the region.

One of the greatest challenges faced has been the design and development of an I-EHR environment to support consistent and authenticated access to clinically significant, multimedia information in order to assist medical decision-making.

In the third version of the I-EHR environment (2001–2003), the main effort was paid on standardization of the indexing mechanism, through an extensive harmonization process that took place as part of the PICNIC project. A patient identification service was introduced, together with the adoption and use of international standards, mainly from the Object Management Group (OMG) in compliance with the PICNIC architecture. A more advanced, multilingual GUI was introduced, and several third-party systems were incorporated into the federation.

Today, the fourth generation of the I-EHR environment is being delivered as a WS, representing its natural evolution. The presented version of I-EHR is mainly based on PICNIC infrastructure work and takes into consideration state-of-the-art developments. It uses W3C technologies to deliver integration across enterprises in a standardized, open, and platform-independent fashion.

A. *Structure of the HYGEIAnet I-EHR*

As pointed out in [13], features with which patients are presented could be designated as problems. By placing the problem list at the front of the clinical record, everyone involved in patient care can be aware of the list of active and inactive/resolved problems.

The domain model applied to the construction of the HYGEIAnet I-EHR is based on the medical encounter that, in its simplest form, consists of a health service, provided by a service provider to a service recipient, to address a health condition at a delivery site at a point or period in time. Therefore, indexing information, maintained by the federated domain model, refers to patient data that are produced during the communication about the patient, between two or more individuals, at least one of whom is a member of the healthcare team currently involved.

Important entities associated with shared encounter information include the patient, the attending healthcare professional, and the clinically significant information the EHR system possesses. Therefore, indexed information is contained within the HII as a list of qualified codes indicating the existence of specific types of clinical information without immediate knowledge of the corresponding actual values.

B. *Componets of the HYGEIAnet I-EHR*

Traditionally, all HII work at the Foundation for Research and Technology–Hellas (FORTH) was based on CORBA technology. In particular, the work of Health Domain Task Force (DTF) of OMG was very important in the definition of the architecture for the I-EHR. This architecture was based on the following services defined by the Health DTF:

1) Person Identification Service (PIDS) for the unique identification of patients;
2) Lexicon Query Service (LQS) for the management of medical terminologies;
3) Clinical Observation Access Service (COAS) for accessing the primary sources of medical information;
4) Resource Access Decision (RAD) Facility for the enforcement of resource-oriented access control policies.

Moreover, in the context of HYGEIAnet, the following services were specified and designed:

1) the I-EHR Indexing Service (I-EHR IS) for managing indexes to the sources of primary information, so that the efficiency and scalability aspects of the architecture are reinforced;
2) the I-EHR Update Broker (I-EHR UB) for keeping the I-EHR IS up to date with new or modified information, and consistent with the information accessible through COAS;
3) the Health Resource Service (HRS) for the unique identification and management of clinical resources, in the context of I-EHR, such as medical stuff, healthcare facilities.

All the above-mentioned services comprise the “backbone” of a component-based computing environment for the provision of the I-EHR service, as shown in Fig. 6.

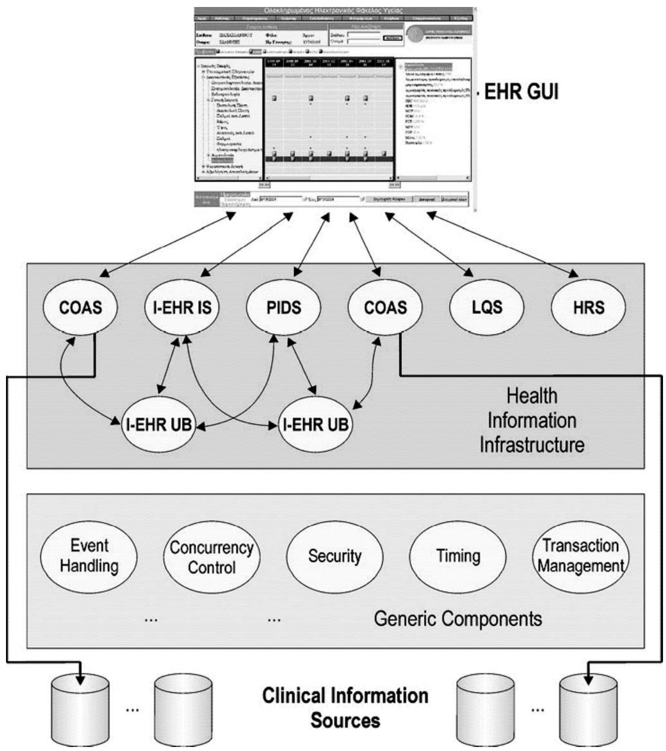

Fig. 6. The HYGEIAnet I-EHR software architecture components consist of COAS, I-EHR IS, PIDS, LQS, HRS, and I-EHR UB. Generic components are provided by platform specific service.

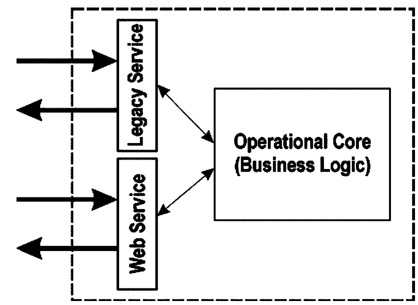

Fig. 7. Enhancing a component's interface with WS functionality.

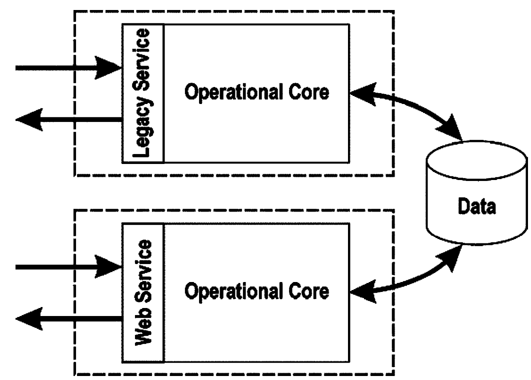

Fig. 8. Implementing a new component with the WS interface.

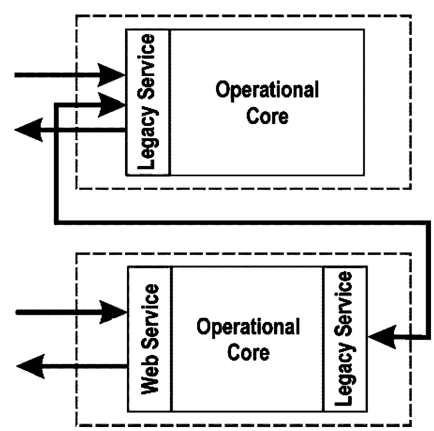

Fig. 9. Implementation of a new WS component that acts as a middleman to the existing one.

C. WS Migration for Cross-Enterprise Integration

The implementation of a WS introduces some interesting challenges, especially in the cases where an existing infrastructure is in place and works. This is because, when migrating from an existing infrastructure toward a new one, several issues with associated costs have to be considered and decisions have to be taken on how and when to upgrade the platform itself and whether the expected benefits and savings outweigh the costs involved.

Exposing a "legacy" component through a WS interface can be performed through one of the following alternative methods.

1) *Incorporation of the WS functionality to the existing component:* This approach requires the enhancement of the interface portion of the component so that it can be accessible through both of its "facades": the legacy (e.g., CORBA) interface and the WS interface (see Fig. 7). If the WS interface was similar to the legacy interface, then it would be desirable to reuse existing code, especially the one related to the business logic. Unfortunately, this is not always feasible, especially if the component's code does not have a clear separation between the interface and the business logic parts.
2) *Implementation of another component offering exactly the same functionality, this time through a WS interface* (see Fig. 8): It requires the implementation of a new WS component having access to the same data used by the legacy component. Now, there is duplication of the business logic and two different source code bases to maintain, something that may cause serious algorithm inconsistency hazards, when offering the new service in parallel with the legacy one.
3) *Implementation of a new WS component that only relays the requests it receives to the already existing component:* In this case, the WS component is just a middleman or a wrapper to the legacy component, which in turn, implements the core functionality (see Fig. 9). What has to be noted is that the WS functionality should be implemented in terms of the existing legacy system interface

and a mapping/translation logic should be in place to transform the WS requests to the legacy service calls and the other way round. This transformation phase has an impact on performance but known techniques like caching and prefetching could be employed in order to improve efficiency.

The implementation of a WS I-EHR presents many challenges since the mapping of an object-oriented programming model and platform (i.e., CORBA), to a service-oriented model and platform is not straightforward. There is ongoing work from OMG to provide a mapping from the CORBA Interface Definition Language (IDL) to the WS Description Language (WSDL) or to define a CORBA binding for WSDL, despite the fact that WS are not distributed objects [14], like those built with CORBA.

D. The I-EHR as a WS

The decision of what should be exposed through a WS interface revolves around what functionalities are expected by the consumers of the service. In the case of the I-EHR, the fundamental driving force is the timely, secure, and accurate provision of the patient's lifetime, clinically significant information to all interested and authorized users.

The I-EHR WS must provide to a consumer the ability to search for a patient's identity given a set of demographic criteria and the retrieval of all the related health and medical information pertaining to the patient under consideration. Additional filtering must be possible if the consumer of the service is only interested in some parts of the patient's medical record.

In the case of HYGEIAnet, the I-EHR WS is built upon the already existing PICNIC architecture; it leverages PIDS for the identification of patients, LQS for the identification of medical terms and the translation of them to descriptions in a natural language, HRS for the identification of resources, I-EHR IS for locating the interesting information, and COAS for retrieving it. The already developed CORBA components are reused and are totally hidden behind the I-EHR WS. This reuse technique permits the exploitation of the existing infrastructure by compliant consumers, so that backward compatibility is assured.

E. Role-Based Access to Information

In adopting a role-based access control (RBAC) approach in a RHIN environment for the support of the I-EHR, all the involved entities, together with their corresponding role, have been identified (see Table II). All HYGEIAnet applications and services are declared within the HYGEIAnet HRS in order to be issued a unique identity. Each user of the I-EHR must also be registered within HRS in order to be eligible for using any of the HYGEIAnet services, and subsequently, to possess a unique username. The user's password is, then, communicated to the authentication server for the facilitation of single-sign-on for all applications and services within HYGEIAnet, while a certificate is also issued by the certification authority. Having completed successfully the registration process, any person with the required credentials becomes a legitimate user of the services provided within the RHIN.

TABLE II
ACTORS INVOLVED IN RBAC

Actor	Role
Health Resource Service (HRS)	It issues unique ID information for users and applications. It maintains and manages public health resource information, activates all applications and services by issuing them a unique ID, registers all healthcare persons and organizations, associates a unique user name for each healthcare persons (i.e. citizen or professional) it maintains, and associates healthcare persons and organizations.
Certification Authority	Issues and maintains certificates for users and applications. The digital certificates are used for digital signing of documents and in the case of users could be stored in smart cards.
Authentication Service	Maintains and manages passwords. It issues passwords for every new healthcare person and performs user authentication.
I-EHR service	Maintains and manages roles (groups) and role-based permissions.

The authorization process, on the other hand, follows a decentralized approach: Each RHIN service defines its own rules that permit or deny access to the authenticated users and applications. During the logon phase, an I-EHR user is authenticated through the centrally controlled authentication server. Subsequently, his/her I-EHR access rights are validated through the I-EHR-specific access control server, which, for the HYGEIAnet implementation, is linked to the I-EHR IS. A new user account in I-EHR can only be created for a user that has already been registered within the RHIN HRS. This includes assignment of roles (and associated permissions) at the IT service level.

Once the RHIN is considered a single security domain, user management is simplified and the single sign on concept can be enforced, while at the same time, various RHIN services can be left free to impose their own access control policies.

At the I-EHR level, the user is assigned to certain roles based on organization (i.e., his/her position in a healthcare provider) or other criteria. The roles are granted permissions and rights that are expressed by "allow" and "deny" rules. Each rule conveys the following information:

1) the type of rule, i.e., if it is allowing or denying access;
2) the source of health information (clinical information system) that this rule applies;
3) the kind of clinical information that this rule applies.

The definition of rules requires the existence of an administrator at the RHIN service level responsible for associating users with roles and assigning rights to the roles. The administrator can create/activate and remove/deactivate users and roles, and associate users with roles.

It is, therefore, feasible to grant access to users based on the location that the information is hosted and based on the type of the information. This model, in spite of its simplicity, has been proved powerful enough for the most common use case scenarios.

Extensions to the definition of the rules that, for example, allow combination of rules and creation of more complex rules using conjunctive and disjunctive constructs, although feasible, have not been needed in practice.

F. Performance Considerations

In application domains where the performance factor is considered to be critical, middleware component developers pay close attention to several performance indicators like throughput, latency, jitter, and scalability. In the case of the HYGEIAnet I-EHR, workload distribution is a natural result to the fact that the original, clinically significant information, is managed locally by autonomous and self-contained distributed data sources at the place produced. It is also a fact that certain HII components, resolving issues related, e.g., to unique person identification, I-EHR indexing, and common terminology, must be controlled centrally.

Services for terminology and health resources, manage information that is not updated frequently and most of the scenarios they participate in are "read" interactions. For such cases, caching techniques are excellent tools for improving performance and fault tolerance. Patient identification and I-EHR indexing services have a mixed workload of read and write operations, so it has been chosen to minimize the state information managed by the service components, and therefore, the backend databases were left to handle the bulk of the data manipulations in short transactions. An example of this design is the implementation of the traversal of a query's results (a "page-by-page iterator"): the results are requested in pieces ("chunks") and every time a new chunk is requested, a new query is executed in order to retrieve only the results of the requested chunk. Although this affects negatively the latency when such a query is performed in isolation, it improves the system's scalability and performance under load since no database locks are held across the network.

By keeping the majority of data in the database tier, horizontal scaling and load balance can be achieved by adding more processing nodes for the service's components that implement the application logic, as long as the database can sustain the load. This has been verified, in the case of PIDS, with datasets of at least four million patients, a number that exceeds, by far, the population of Crete. When larger datasets or throughput should be managed, database technologies like *replication* or techniques like *horizontal partitioning* of data could be employed so that the database tier keeps up with the load.

In the case of the WS component that is located between the I-EHR infrastructure and the clients, the situation is simpler since no data is managed. Here, a cluster of Web servers is able to sustain throughput, as long as the I-EHR infrastructure permits it, and caching/prefetching methods further increase performance. Of course, the use of XML as the marshalling format of data results in less compact messages, and therefore, demands more bandwidth in order to be transmitted and more processing power in order to be parsed than do certain binary communication protocols such as CORBA's Internet Inter-ORB Protocol. This is more evident when binary data should be transmitted: since XML represents textual data, binary information should be encoded in a text representation using, for example, the Base64 encoding algorithm. These transformations increase the processing time needed for the serialization of the data and also the size of the messages, since Base64 encoded messages are generally 33% larger than the original messages. The W3C has proposed the SOAP Message Transmission Optimization Mechanism and the XML-binary Optimized Packaging for treating binary data efficiently inside SOAP and XML messages in general. Furthermore, alternative ways to transmit preparsed XML messages in a binary format are being considered [15], [16], so that a receiving application need not have the overhead of parsing the XML stream.

VI. Discussion

Because of the fact that today, the delivery of a lifelong I-EHR for sharing patient information is considered a crucial factor in reducing medical errors, our ability to link unaffiliated sites to share patient data with each other (requiring the communication of complex and diverse forms of information between a variety of clinical and other settings) is expected to close the information gap that has traditionally impaired the delivery of the highest quality of care.

A lifelong I-EHR has been piloted in HYGEIAnet in Crete, Greece, for the past few years, under alternative operational settings. In all cases, effort was paid to sustain the autonomy of the local clinical information systems managing clinically significant information at the place originally produced. A number of clinical and operational benefits have been revealed [17]. At the same time, user confidence, in relation to the trust and security infrastructure of the RHIN, has been proven to be an important factor influencing user adoption.

A major problem related to the actual deployment of an I-EHR service, as with any other form of technology, is that no matter how advanced a system may be, the users are the ones who will have the final word, accepting or discarding it. The response of physicians is shaped by their perception of how new practices affect existing relations and the accomplishment of professional goals. Legal and standards-related barriers were also encountered, relating to privacy, confidentiality, and security, as well as the lack of standards for healthcare information systems.

The technical issues, as described in previous sections, although of considerable depth and spanning, have not been a significant barrier for the delivery of a lifelong I-EHR service. A large number of re-usable software components (middleware services) have been identified, developed, deployed, and piloted, capable of supporting the ICT "evolutionary and migration strategy" at the regional level. Our experience has shown that re-usable components make the deployment of new clinical information systems and e-health services easier and faster, while at the same time, they allow for a more efficient management of the ICT infrastructure. Still, decisions have to be formulated about how and when to upgrade the platform itself and whether the expected benefits and savings outweigh the costs involved.

It is, indeed, a fact that the idea of exposing some existing functionality through some new interface is not new, and in fact, it reappears every time a new technology emerges. Clearly, migration toward a new infrastructure should be planned carefully, and in cases where significant investments are required, the implementation of new interfaces using already existing infrastructure is more appropriate, especially if it is a proven one and it builds on existing human effort and experience. This is more than true for WS since they push interoperability in a higher layer: one where different middleware technologies need to interoperate.

ACKNOWLEDGMENT

The authors would like to thank all members of the Biomedical Informatics Laboratory of ICS-FORTH who had been active in and out of the various I-EHR and HII projects since 1995. They also acknowledge Prof. S. Orphanoudakis, Founder and Leader of the Laboratory, and elected Director of FORTH, until his premature death in March 2005.

REFERENCES

[1] M. Tsiknakis, D. G. Katehakis, and S. C. Orphanoudakis. (2002, Dec.18). An open, component-based information infrastructure for integrated health information networks. *Int. J. Med. Inf.* [Online] *68(1–3)*, pp. 3–26. Available: http://www.sciencedirect.com/science/article/B6T7S-47BXBTX-2/1/ca92f2cd3c77a1f7be7073beef4bf493

[2] G. T. W. Dietzel and C. Riepe, "Modernizing healthcare in Germany by introducing the eHealthcard," *Swiss Med. Inf.*, no. 52, pp. 18–22, 2004.

[3] ISO/TC 215*, Health Informatics, Electronic Health Record Definition, Scope, and Context,*, Standard ISO/TR 20514:2005, Oct. 17, 2005.

[4] ISO/TC 215*, Health Informatics, Requirements for an Electronic Health Record Reference Architecture*, Standard ISO/TS 18308:2004, Jan. 22, 2004.

[5] *Health Informatics—Electronic Healthcare Record Communication—Part 1: Extended Architecture*, Standard CEN/TC 251 ENV 13606-1:1999.

[6] *Health Informatics—Electronic Healthcare Record Communication—Part 2: Domain Termlist*, Standard CEN/TC 251 ENV 13606-2:2000.

[7] *Health Informatics—Electronic Healthcare Record Communication—Part 3: Distribution Rules*, Standard CEN/TC 251 ENV 13606-3:2000.

[8] *Health Informatics—Electronic Healthcare Record Communication—Part 4: Messages for the Exchange of Information*, Standard CEN/TC 251 ENV 13606-4:1999.

[9] N. Saranummi, D. Piggott, D. G. Katehakis, M. Tsiknakis, and K. Bernstein, Eds., *Regional Health Economies and ICT Services—The PICNIC Experience, vol. 155, Studies in Health Technology and Informatics*. Amsterdam, The Netherlands: IOS Press, Aug. 2005, 248 pp. (ISBN: 1-58603-538-x).

[10] *IEEE Recommended Practice for Architectural Description for Software-Intensive Systems*, IEEE Standard 1471-2000.

[11] ISO/TC 215*, Health Informatics, Trusted End-to-End Information Flows*, Standard ISO/TR 21089:2004, Jun. 15.

[12] D. G. Katehakis, M. Tsiknakis, and S. C. Orphanoudakis, Enabling components of HYGEIAnet. presented at TEPR 2001, Boston, MA, May. 8–13, pp. 146–153, [Online]. Available: http://www.ics.forth.gr/˜katehaki/publications/tepr2001.pdf

[13] L. L. Weed, *Medical Records, Medical Education, and Patient Care—The Problem-Oriented Record as a Basic Tool*. Cleveland, OH: Case Western Reserve Univ. Press, 1969.

[14] W. Vogelsm, "Web Services are not distributed objects," *IEEE Internet Comput.*, vol. 7, no. 6, pp. 59–66, Nov./Dec. 2003.

[15] ITU-T Rec. X.694 | ISO/IEC 8825-5. (2004). *Information Technology—ASN.1 Encoding Rules: Mapping W3C XML Schema Definitions into ASN.1* [Online]. Available: http://www.itu.int/ITU-T/studygroups/com17/languages/X.694.pdf

[16] *Information Technology—Generic Applications of ASN.1: Fast Web Services*, ITU-T Rec. X.892|ISO/IEC 24824-2, 2005.

[17] M. Tsiknakis, A. Kouroubali, D. Vourvahakis, and S. C. Orphanoudakis, "Implementing a Regional Health Information Network: Impact on health care performance and the management of change," *Int. Health Care Manage. Adv. Health Care Manage.*, vol. 5, pp. 297–329, 2005.

Dimitrios G. Katehakis (S'88–M'93) was born in Heraklion, Crete, Greece, in 1968. He received the Diploma from the Technical University of Patras, Patras, Greece, in 1991, and the M.S. degree from the University of Maryland, College Park, in 1993, both in electrical engineering.

He was with the Medical School, University of Maryland, Baltimore, for Project IONDT (U.S. Federal) and Project MCR (MD State, Department of Health and Mental Hygiene, CDC). He was also with the National Defense Ministry of Greece and the Hellenic Telecommunications Organization (OTE). He is currently a Senior Telecommunications Engineer at the Biomedical Informatics Laboratory, Institute of Computer Science, Foundation for Research and Technology–Hellas, Heraklion, Greece. He has participated in several European RTD projects including RETRANSPLANT (DGXIII—Health Telematics), TEMeTeN (DGXVI—RISI 2), HECTOR (DGXIII—Health Telematics), and PICNIC (IST-1999-10345). His current research interests include data integration analysis and the development of the appropriate Healthcare Information Infrastructure (HII) required for the provision of advanced health information network services.

Mr. Katehakis is a member of the Association for Computing Machinery (ACM) and the Technical Chamber of Greece.

Stelios G. Sfakianakis was born in Heraklion, Greece, in 1974. He received the B.Sc. degree in computer science and the M.Sc. degree in advanced information systems from the University of Athens, Athens, Greece, in 1995 and 1998, respectively.

Since January 2000, he has been with the Biomedical Informatics Laboratory, Institute of Computer Science, Foundation for Research and Technology–Hellas, Heraklion, where he is engaged in the design and implementation of a component-based architecture for the Integrated Electronic Health Record using CORBA and Web Services as the middleware technology. His current research interests include the design of relational or hierarchical databases to visual modeling with UML, open source operating systems, and distributed object and service-oriented computing.

Georgios Kavlentakis was born in Heraklion, Greece, in 1971. He received the Electrical and Computer Engineering Diploma in computer and electronics from Aristotle University of Thessaloniki, Thessaloniki, Greece, in 1996.

Since April 1999, he has been with the Biomedical Informatics Laboratory, Institute of Computer Science, Foundation for Research and Technology–Hellas, Heraklion, where he is engaged in the implementation of an Integrated Electronic Healthcare Record using CORBA as the middleware technology. He has also developed Primary Health Care Information System using ASP technology. His current research interests include the design of relational or hierarchical databases, distributed computing and programming using C++, Visual Basic, PowerBuilder, and the design and implementation of Web-based applications.

Dimitrios N. Anthoulakis received the M.Sc. degree in interactive multimedia at Heriot-Watt University, Edinburgh, U.K., in 2001.

From September 2001 to March 2006, he was with the Biomedical Informatics Laboratory, Institute of Computer Science, Foundation for Research and Technology–Hellas, Heraklion, Greece, where he was engaged in the development of HYGEIAnet, which is the Regional Health Informatics Network of Crete. Since April 2006, he works with Usableweb Ltd., Heraklion, Greece. He has designed and developed both a Java and a .Net web client for the Integrated Electronic Health Record as a deliverable of PICNIC European project, using open middleware technologies based on CORBA and Web services. His current research interests include information retrieval and a user-centered interface design approach in Health Care Environments, HCI usability issues, and evaluation and spatial representation of an integrated electronic healthcare record.

Manolis Tsiknakis (M'92) received the B.Eng. degree in electronic engineering, in 1983, the M.Sc. degree in microprocessor engineering, in 1985, and the Ph.D. degree in control systems engineering from the University of Bradford, Bradford, U.K., in 1989.

Since 1992, he has been with the Institute of Computer Science, Foundation for Research and Technology–Hellas, Heraklion, Greece, where he is currently a Principal Researcher in many collaborative R&D projects, the Head of the Center of eHealth Technologies, and is engaged in the development of HYGEIAnet, the integrated health information network of Crete. He is the initiator and Co-Chair of the ERCIM Biomedical Informatics Working Group. His current research interests include biomedical informatics, component-based software engineering, information integration, ambient intelligence in eHealth and mHealth service platforms, and signal processing and analysis.

Dr. Tsiknakis was a member of the Programme Committee of the HealthGRID from 2003 to 2005 and the IEEE CBMS 2005 International Conferences. He is currently a member of the Management Board of HL7–Hellas and the Association for Computing Machinery (ACM), and a National Representative in the EUMEDIS programme of the EU.