WebOnCOLL: Medical Collaboration in Regional Healthcare Networks

Catherine E. Chronaki, Dimitrios G. Katehakis, *Member, IEEE,* Xenophon C. Zabulis, Manolis Tsiknakis, *Member, IEEE,* and Stelios C. Orphanoudakis, *Senior Member, IEEE*

***Abstract*—This paper presents *WebOnCOLL*, a web-based medical collaboration environment, which has been designed in the context of the regional healthcare network of Crete. WebOnCOLL employs the infrastructure of regional healthcare networks to provide integrated services for virtual workspaces, annotations, e-mail, and on-line collaboration. Virtual workspaces support collaborative concepts like personal web pages, bulletin boards, discussion lists, shared workspaces, and medical case folders. Annotations provide a natural way for people to interact with multimedia content, while e-mail is one of the most popular forms of communication today. On-line collaboration satisfies the need for a more direct form of communication.**

***Index Terms*— Computer-supported cooperative work (CSCW), healthcare information infrastructure, medical collaboration, medical digital libraries, regional healthcare networks, shared workspaces.**

I. Introduction

A REGIONAL network provides healthcare professionals with the necessary infrastructure to collaborate with their peers, share opinions, exchange clinical data, and access domain-specific information within or across regions. *WebOnCOLL* is a web-based environment that allows healthcare personnel to collaborate with their peers regardless of geographic location. WebOnCOLL employs the healthcare information infrastructure of the evolving regional healthcare network of Crete, Greece, which includes high-speed networks and regional directory, resource, and security services that are being developed as enabling middleware components to provide integrated web-based collaboration services for virtual workspaces, annotations, e-mail, and on-line collaboration.

This paper describes the architecture and services of the WebOnCOLL collaboration environment. Central notions of the WebOnCOLL architecture are user profiles and virtual workspaces. User profiles allow the customization of the environment to user authorities, tasks, and preferences. Each user session is associated with a heterogeneous data collection: a virtual workspace. A virtual workspace maintains session information, service results, and multimedia objects that the user selects or creates during navigation. Virtual workspaces are public, private, or shared and persist across services and user sessions. Thus, in addition to session management, virtual workspaces provide for secure collaboration, data availability, and persistence. In WebOnCOLL, virtual workspaces support a number of collaborative concepts, i.e., personal web pages, bulletin boards, discussion lists, shared workspaces, and medical case folders. Medical case folders maintain information relevant to a medical case, e.g., images, examinations, comments, discussions, updates, alerts, etc.

The rest of the paper is organized as follows. Section II provides background on web-based collaboration. Section III discusses the integrated services of the WebOnCOLL collaboration environment. Section IV presents the architecture of WebOnCOLL servers. The discussion focuses on the role of the workspace and user profile managers. Then, Section V presents future work, and Section VI raises usability and accessibility issues. Finally, Section VII concludes the paper.

II. Background

Although the web was originally conceived as a means of information sharing and dissemination, its network and operating system transparency promote its use as a collaboration platform for geographically distributed working groups. In fact, the World Wide Web (WWW) is collaborative technology in the weak sense of the word, since it allows people to share information. Several issues are related to the development of web-based collaborative applications. These include the architecture, information structure, availability, security, awareness, and compatibility [1]. Core functional components of any collaboration environment are *content*, *communication*, and *control* [2], [3]. Content refers to the domain objects that the users collaborate on, such as shared documents, pointers, and hyperlinks. Normally, the unit of information on the web is the page. However, systems developed for computer-supported cooperative work (CSCW) adhere to directory/folder hierarchies treated as conferences in bulletin boards. Some applications have a more explicit model of the web network structure, while others recognize portions of pages. Availability and replication of content are issues that need to be addressed by web-based collaboration services. Hyperlinks (URL's) offer an easy way to refer to meta-information about content and relevant services. The emerging HTML 4.0 and XML standards address several issues related to content structure, availability, and replication [4], [5].

Manuscript received November 10, 1997; revised December 23, 1997.

C. E. Chronaki, D. G. Katehakis, X. C. Zabulis, and M. Tsiknakis are with the Center for Medical Informatics and Health Telematics Applications, Institute of Computer Science, Foundation for Research and Technology, Heraklion, 71110 Crete, Greece (e-mail: chronaki@ics.forth.gr).

S. C. Orphanoudakis is with the Computer Science Department, University of Crete, Heraklion, 71110 Crete, Greece, and the Center for Medical Informatics and Health Telematics Applications, Institute of Computer Science, Foundation for Research and Technology, Heraklion, 71110 Crete, Greece.

Publisher Item Identifier S 1089-7771(97)09417-X.

Communication may be synchronous or asynchronous. As pointed out in [1], the bulletin board style of interaction that transforms communication into browsing an information structure is the type of interaction that the web was designed to manage. Indeed, the nature of the web involves local update and remote browsing. The desire for more symmetric remote access leads inevitably toward issues of concurrent update, locking, access control, and version management. In addition, the need for high availability increases the use of caches and complicates the replication process.

In cooperative work, people interact with objects to control their behavior and update their state. Feedback, feedthrough, notification, and awareness are import aspects of control [20], [39]. Feedback concerns a user's perception of his/her own actions, while feedthrough concerns a user's perception of the effects of other people's actions. Notification concerns a user's perception of occurring events, while awareness refers to perception of the presence of other people. In a web-based collaboration environment, feedback is best accomplished by downloading the relevant objects before acting upon them. However, this reduces feedthrough for the collaborators. Feedthrough, awareness, and notification are not supported directly by the web infrastructure and are frequently provided by special protocols.

The web infrastructure comprises server software (web servers), an application protocol (HTTP), data format (HTML), and client software (web browsers). Various web-based collaboration applications modify or extend parts of this infrastructure in an attempt to overcome the inherent limitations for peer-to-peer communication [6]–[8]. Such extensions or modifications include server customization, "helper" applications and applets that promote client-side computing, and special protocols for rapid interaction or notification.

Other important architectural issues concern information about people as well as security and confidentiality services. The requirements of the target user group should be also taken into account. For example, a document repository for public use has different requirements from one installed in a company intranet. Automated agents and traders further complicate the picture.

WebOnCOLL targets medical collaboration across regional networks. In this context, security is a major issue, which will be addressed on a case-by-case basis, orthogonally to the services provided by the collaboration environment.

A. Collaboration Infrastructure

The collaboration services of a web-based environment include application sharing, audio/video conferencing, bulletin boards, directory services, e-mail, fax, file transfer, group scheduling, Internet fax, Internet phone, listservers, text-based chat, voice mail, webcasting, and whiteboards. These tools facilitate interaction, depending on the static or dynamic nature of information exchange, the synchronous or asynchronous nature of the session, and the computer system configuration of the collaborators [9], [10], [38]. Fig. 1 displays a two-dimensional (2-D) classification of these tools according to the nature of the services they offer. Synchronous and asynchronous collaboration differ in several points; the grain size of action, how the participants come to see the most recent actions, the time frame in which actions become available to other participants, and whether actions can be reviewed [11]. The current technology trend pushes functionality toward the bottom two quarters of Fig. 1, increasing the mobility of participants and the nondeterminism in the interaction. In this trend, locating the participants is a limiting factor addressed by locator and directory services [12], [13]. It is worth noting that this collaboration infrastructure is available in fourth generation browsers, such as Netscape Communicator and Microsoft Internet Explorer [16].

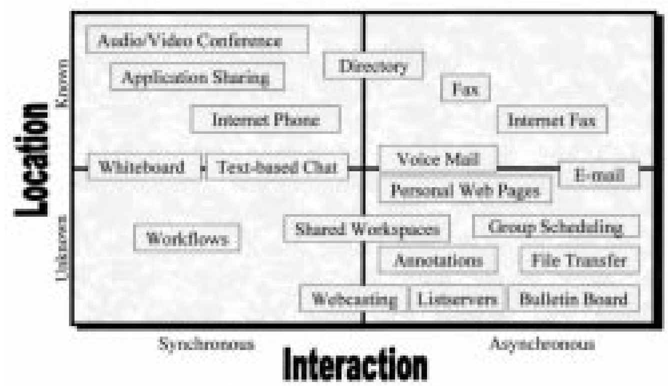

Fig. 1. Classification of web-based collaboration services according to their time and space requirements.

Audio/video conferencing facilitates synchronous group collaboration. Already, low-bandwidth compatibility and standardization efforts (e.g., the ITU H.320 standards for video conferencing [14]) enable interoperability with products from various vendors. Provided adequate bandwidth is available, *video conferencing* transmits and receives, over the Internet, video data during conferences between two or more participants at different sites. A *whiteboard* allows multiple users to display images and draw on a common area that may use multiple screens at various resolutions. *Text-based chat* is frequently a companion functionality to that of a whiteboard in real-time on-line communication. *Applications sharing* facilitates the use of an application collaboratively, e.g., editing a lab report or viewing an X-ray or an ECG, even if the computer of a participant doesn't have that application installed. The term "collaborative browsing" is used for application sharing when the application resides on only one of the machines connected to the conference.

E-mail is the single most successful application for the asynchronous exchange of electronic messages. *Internet phone* performs functions traditionally performed by telephone equipment (e.g., answering machine). *Voice mail* attaches compressed audio files to standard e-mail. *Shared information spaces*, *annotations*, *listservers*, and *bulletin boards* facilitate activities, which involve collaborative posting and reviewing of electronic messages or documents. *Group scheduling* allows colleagues to view each other's calendar and thus help them schedule meetings and appointments. *Fax* and *file transfer* (multipoint or batch) are different methods of document sharing. *Internet fax* is used for the routing of e-mails (or faxes) to faxes through the Internet instead of traditional telephone networks.

Directories make the physical network topology and protocols transparent, facilitating access to people and their preferences. General or domain-specific yellow pages, which are based on LDAP and X.500 [12], [13], are becoming increasingly available.

Webcasting and information channels employ push technology to enable the receipt of information without requesting it [48].

The primary design objective of WebOnCOLL is to provide an integrated set of collaborative services in the context of an open collaborative environment customized for the healthcare domain. Consider, for example, a network of experts that provide healthcare-related consultation services. Within the network, private conversation, secure e-mail, audio and video conferencing, as well as personal web pages facilitate the exchange of information. In addition, information that is generated for the purpose of furthering the profession needs to be accessible through appropriate channels like journals, websites, conferences, listservers, and bulletin boards.

B. Related Work

Information sharing, communication, and coordination are basic elements of any collaborative endeavor [17]. Recently, the notion of a collaboratory, a "center without walls," in which researchers can perform their research without regard to geography has received a lot of attention [18]. CREW [19] is a collaboratory supporting a community of researchers located at The University of Michigan, Ann Arbor, and sponsoring corporations. The CREW desktop is a computer-based environment, which provides members of the CREW community with access to other members, information about CREW-related events, and shared documents. Parts of the CREW desktop provide subscribers with information about other CREW researchers, their activities, location, availability, and interests as well as access to general events calendars, individual calendars, activity in the CREW desktop, and activity on their workstations. In addition, collaboration tools like postable notes and e-mail are also available [19].

In the healthcare domain, researchers from hospital and academic institutions are participating in the InterMed collaboratory, which targets web-based collaboration in medical informatics [21]. The principal goals of this project include promoting the development of healthcare information systems as a collaborative activity, developing a robust framework for collaboration, and using the Internet as the primary vehicle for such collaboration. Specific tasks of the project include security and confidentiality services, vocabulary, access to on-line knowledge, clinical data access and sharing, as well as collaborative policies, procedures, and guidelines.

Several telemedicine projects approach healthcare delivery as a collaborative activity. In the TeleMed project [22], the virtual patient record forms the basis for a collaboratory environment in which multiple physicians and ultimately the patient should be able to engage in an interactive electronic discussion. The TeleMed architecture enables both real-time consultations and off-line sessions. In particular, TeleMed supports real-time interactive collaboration between multiple users. Multiple physicians at remote locations can simultaneously view, edit, and annotate patient data. Each physician can see the data another physician has entered as well as monitor some of the other physician's interactions.

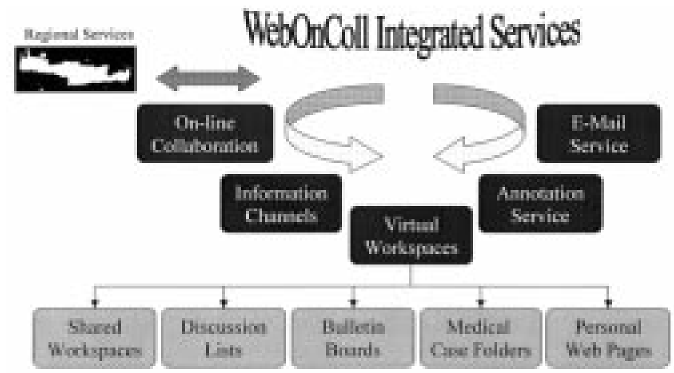

Fig. 2. Integrated collaboration services in WebOnCOLL.

The ARTEMIS project aims to "advance cooperative activities of healthcare providers to promote the delivery of total and real-time care [17]." A multidisciplinary group of computer systems developers, physicians, and healthcare researchers are using research prototypes and commercial off-the-shelf technologies to develop an open collaboration environment for the healthcare domain. In ARTEMIS, community care networks consisting of a collection of primary care and specialized care providers collaborate to meet the healthcare needs of the community. This communication infrastructure allows primary care physicians to consult with remote specialists in the areas of perinatology and radiology, with computer support for X-rays, ultrasound, voice-annotations, and other multimedia information. The requirements of WebOnCOLL are similar to those of ARTEMIS in that both environments will employ research prototypes and off-the-self components to promote collaboration in healthcare.

Another facet of collaboration that enables the citizens to make informed decisions about their health, relates to the concept of the virtual medical office. The virtual telemedicine office [23] is an integrated environment that encourages users to take active participation in the management of their health. It provides access to digital medical libraries, yellow pages, and regional healthcare resources. Several of the services provided by the virtual telemedicine office, such as directory services, access to digital medical libraries, and teleconsultation tools will also be provided by WebOnCOLL.

III. WebOnCOLL Integrated Collaboration Services

A. Overview

A web browser provides access to the integrated services of WebOnCOLL, as shown in Fig. 2, which facilitate asynchronous as well as synchronous collaboration. E-mail, annotations, and the various forms of virtual workspaces, (i.e., personal web pages, bulletin boards and discussion lists, shared workspaces, and medical case folders) are asynchronous collaboration media. The on-line collaboration service of WebOnCOLL or a third-party conference software like Microsoft NetMeeting or Netscape Conference [16] provides synchronous collaboration services. The maintenance of session information and a common context by virtual workspaces provides service interoperability, while user profiles customize the environment in accordance to user preferences. Thus, virtual workspaces facilitate collaborative activities, while user profiles provide personalized collaboration services. The following sections (Section III-B–E) discuss the integrated collaboration services that will be provided in WebOnCOLL. Section III-F discusses regional information services, as WebOnCOLL employs them to achieve a seamless integration of its services with the regional healthcare infrastructure.

B. Virtual Workspaces

Virtual workspace is a term that originated in the virtual reality (VR) community to denote a three-dimensional (3-D) workspace whose contents can be viewed and examined [26] although it does not actually exist in the real world. In the context of I^2*Cnet* [32], a network of image servers with content-based search capabilities, virtual workspaces were employed to provide for service interoperability, user collaboration, and session management. In this sense, the concept of virtual workspaces resembles "collectors," as employed in the grassroots system [25]. Grassroots collectors are generic elements, which represent collections, such as e-mail folders (collections of mail messages), newsgroups (articles), web pages (HTML links), mailing lists (addresses), access-control groups (identities), etc. Collectors have inflow and outflow, and there are a few canonical operations that can be applied to them. In this view, collaboration between people is mediated by collections; people communicate with each other by transferring information between their "collectors." The main difference is that virtual workspaces in WebOnCOLL are heterogeneous collectors, and thus, it is the workspace objects that specify the acceptable operations.

Virtual workspaces maintain session information, history of interaction, and access privileges at the workspace level. Hence, they provide the context in a user session. The user may populate the current virtual workspace with objects from an intranet, the web, or private collections. The workspace maintains a "contents" file that contains the location (URL) of each workspace object. The only case when a workspace object physically exists within the WebOnCOLL server is when the workspace object is the result of a service request. The results of a service request are automatically inserted in the workspace. An interesting variation of the workspace concept is the information channel. Information channels employ push technology to update their contents, based on user preferences, as recorded in the user profile [48].

Virtual workspaces may be public, shared, or private. Public workspaces are used to receive comments, postings, or annotations. Shared workspaces are used to provide the collaboration material in a synchronous or asynchronous collaboration session. Private workspaces are personal web pages, which are protected and persist in time. The owner of a workspace may selectively share objects of the workspace with collaborators by adjusting the access control information of individual workspace objects. Virtual workspaces are hierarchical, in the sense that a virtual workspace may contain other workspaces.

In WebOnCOLL, virtual workspaces currently support personal web pages, bulletin boards and discussion lists, shared workspaces, and medical case folders.

1) Personal Web Pages: In WebOnCOLL, personal web pages are workspaces that serve both as a contact point and as a front-end to the working environment of the user. A personal web page is a contact point because people that wish to communicate with the user may post notes or content to a personal bulletin board. At the same time, the personal web page provides access to the working environment of the user. The working environment includes pointers to private and shared workspaces, information channels, and other WebOnCOLL services. Personal web pages form the front-end of the virtual medical office, which is part of our future work (see Section V). Collaborators or patients may post second-opinion requests as medical case folders. The notification support of virtual workspaces ensures that the physician is informed of the updates.

2) Bulletin Boards and Discussion Lists: Traditionally, bulletin boards and discussion lists have been the grounds for active social interaction. Thus, authorized users may collaborate through postings to appropriate regional bulletin boards or discussion lists. They may post their comments, views, opinions, or propose revisions and updates on the contents of an information repository. In WebOnCOLL, bulletin boards are special forms of virtual workspaces. Bulletin boards may be associated with closed collaboration groups, as in the case of researchers collaborating on health monitoring, which is also part of our future work (Section V).

3) Shared Workspaces: The concept of shared information spaces plays a central role in web-based collaborative activities [6]–[8], [27], irrespective of the nature of the interaction (asynchronous or synchronous). The Basic Support for Cooperative Work (BSCW) system [6] is a shared document management system, organized around the notion of collaboration groups, which may jointly author, comment, and annotate documents. A BSCW workspace may contain various forms of electronic objects, such as documents, spreadsheets, and pictures or links to WWW pages, and provides basic support for awareness, authentication, authorization, and version control. DreSS [8] is a multiuser web-publishing environment based on a combination of forms, CGI-scripts and two small client-side programs. When compared to BSCW, DreSS is much simpler and does not support awareness. Hyper-G [27] is a hypermedia system, which supports powerful browsing and link-management. Hyper-G can be accessed through web browsers and supports distributed web authoring. The Futplex system [7] provides shared documents with read/write access. A Futplex document consists of a tree structure of web pages. Pages can be created, changed, and added interactively by using a standard web client.

Shared workspaces in WebOnCOLL are not concerned with issues like joint authoring of documents and version control. In contrast, important issues are information organization and

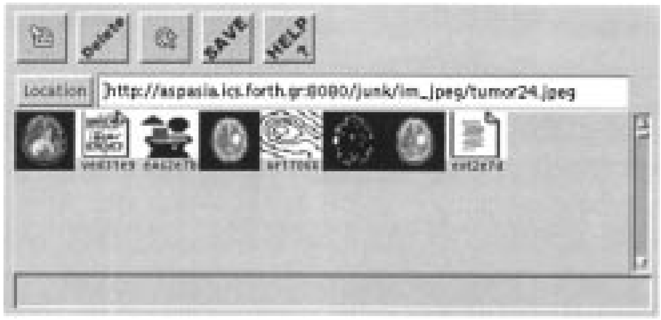

Fig. 3. Shared workspaces provide transparent access to multimedia data (java applet interface to a virtual workspace).

availability, session management, awareness, and notification. If the functionality provided by the BSCW system [6] is considered necessary, BSCW workspaces will be provided in WebOnCOLL as an added-value service. Shared workspaces in WebOnCOLL are hierarchical in that they may contain other workspaces, and support user cooperation by providing secure and transparent access to a shared data collection. This data collection includes multimedia data objects (images, voice, video, patient exams, annotations, etc.) inserted by the users sharing the workspace (Fig. 3). Thus, multiple users collaborate over a workspace, sharing material of common interest. All users connected to a shared workspace are notified of workspace updates as they occur and may inspect the properties of workspace objects, which include author, last update, and a brief description. Furthermore, a user connected to a workspace may be aware of other users also connected to the workspace. Shared workspaces are also used to organize the conference material for an on-line collaboration session. Hence, conference management in WebOnCOLL follows a shared information space approach, according to which, the conference initiator simply provides the name of a shared workspace. No further actions are necessary to provide the collaboration material to conference participants.

4) Medical Case Folders: The abstraction of the medical case folder, combined with the provision of "helper" applications and web services for viewing, processing, and collaborating on its multimedia contents, sum up the approach to medical collaboration, as employed in WebOnCOLL. A medical case folder simplifies the organization of multimedia information related to a particular encounter. Collections of medical case folders facilitate various forms of collaboration, such as teleconsultation, referral, and education.

A typical medical case folder includes information from medical history, medication, progress reports, and data-intensive diagnostic services, such as laboratories, imaging departments, etc., following the Subjective, Objective, Assessment Plan (SOAP) information model [44]. Part of the multimedia data in a typical medical case folder requires specialized applications for viewing and processing, e.g., DICOM or ECG viewers. To meet this requirement, WebOnCOLL provides "helper" applications, which are associated with medical data types.

A medical case folder is a virtual workspace. The owner of the workspace (medical case folder) is able to add patient record data either through direct access to web-enabled clinical information systems or through the virtual patient record interface, which integrates and provides access to the heterogeneous autonomous information systems within the region. A medical case folder may be linked to information channels, which provide customized access to diagnostic databases and digital medical libraries. Furthermore, access to regional information services facilitates secure and up-to-date information.

The medical case folder provides a partial view of the electronic patient record. An active inclusion of patient-related information in the medical case folder has to take into account aspects of presentation, control, and functional integration. Presentation integration is partially facilitated through the use of the web as the user interface platform. Control integration is required to securely access clinical information systems to retrieve additional patient information. Currently, control integration is possible only for web-enabled information systems. Functional integration is required for the workflow-based execution of healthcare processes, such as examination ordering and booking [57]. Elements of a domain-specific framework for the integration of distributed healthcare record segments, which form the logical components of the integrated healthcare record [28], will be combined with regional booking services [29] to meet these functional integration requirements.

C. E-mail

E-mail, one of the most popular forms of communication today, has led to the emergence of on-line communities by allowing people in different geographical areas to communicate asynchronously. Presently, e-mail has evolved to a point where it is now used for multiple purposes: *document delivery and archiving*, *work task delegation*, and *task tracking*. Thus, e-mail is becoming the single most important form of communication, serving as the source of many office tasks [43]. In WebOnCOLL, healthcare professionals may mail to each other any data object present in a virtual workspace using the built-in e-mail support of any advanced web browser. Furthermore, the combination of user profiles and the advanced features of fourth generation browsers ensure that mail is accessible and manageable from any place and at any time.

D. Annotation Service

The success of the web has triggered significant research activity on web-based annotation systems [15], [40], [41], [49]. This is partly due to the fact that annotations provide a natural way for people to interact with multimedia content. The ComMentor system, which has been developed in the context of the Stanford Digital Library Project, is a powerful annotation environment that allows people to share structured in-place annotations for arbitrary documents [40], [41], [49]. Based on meta-information, ComMentor selects the annotations that should be merged with the document, creates an annotated document, and delivers it to the browser.

The annotation service of WebOnCOLL, which was originally developed in the context of I^2*Cnet* [31], complements the work of systems like ComMentor, in that it applies to

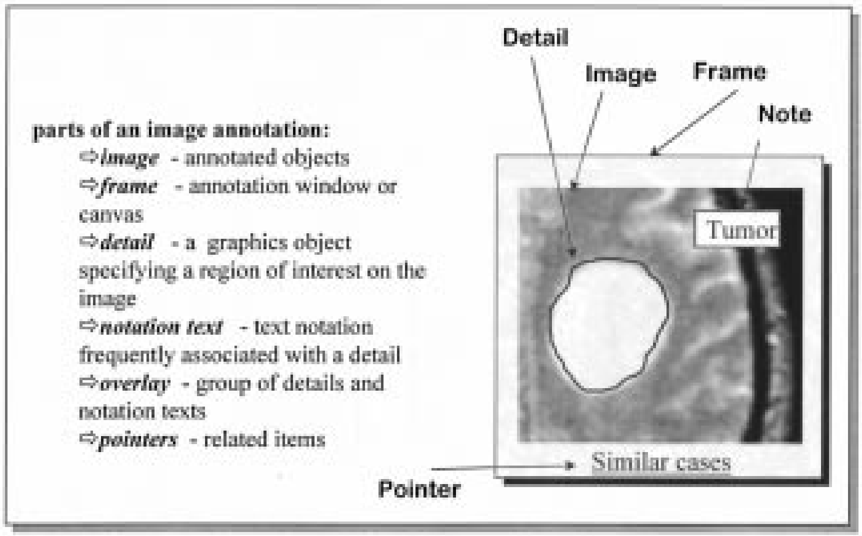

Fig. 4. The native annotation format promotes the expression of different forms of annotation: comment, refutation, confirmation, correction, and illustration.

groups of images, rather than general-purpose documents or single images. In this way, the user may compare multiple images or comment on a trend represented by a sequence of images. Another aspect of the annotation service is that each annotation is associated with *pointers* indicating relevant digital objects. Pointers are uniform resource locators (URL's), which point to a multimedia document, an audio file, or another annotation. As shown in Fig. 4, an annotation includes one or more *details* that are graphical objects specifying a region of interest in an annotated image. Frequently, a detail is linked to *notation text*. Notation text is a textual field, which may refer to the region of interest outlined by the detail or the annotated image as a whole. Multiple details and notation texts may be grouped together to form *overlays*, which can be viewed concurrently or one at a time. Additionally, an *annotation message* enables the inclusion of a text report or message that supplements the graphical part and the pointers of the annotation. This message may include an overview of the purpose of the annotation. Finally, each annotation is associated with a *property* object, which maintains meta-information, including the author of the annotation, creation date, review date, moderator, subject, and relevant keywords. Meta-information is used for annotation indexing and search.

The current user interface of the annotation service, as shown in Fig. 5, 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 fora and on-line collaboration sessions.

E. On-Line Collaboration

WebOnCOLL users have the option of using commercial or third-party products for synchronous and asynchronous collaboration and interaction. By using an off-the-shelf customizable commercial application (e.g., NetMeeting [16], CU-SeeMe [37], and CoMed [53]), medical professionals may share and collectively control viewer applications, e.g., a DICOM viewer or a specialized image-processing environment (DIPE [33]). At the same time, WebOnCOLL provides tools for organizing the collaboration material in medical case folders, sharable viewers for multimedia data, as well as access to regional information services.

The on-line collaboration service of WebOnCOLL is a lightweight Java environment, which provides shared white-board, text chat-box, and access to a shared workspace. Thus, it satisfies the need for a lightweight, but direct form of communication. Users sharing a workspace may exchange opinions and ideas by using a text-based chat facility and may also comment on various workspace objects. In the course of an on-line session, the joint annotation of diagnostic images in a medical case folder can be performed.

When an authorized user requests the creation of a collaboration session bound to a specific workspace, a new instance of the collaboration server is launched. The contents of the workspace at that time constitute the conference material, and the user that requests the launching of the collaboration session controls the floor. All messages pass through the collaboration server, enabling the recording of the complete collaboration session. Furthermore, snapshots of the session may be stored in the annotation format.

F. Regional Information Services

The infrastructure of a regional healthcare network comprises high-speed networks and regional information services. Regional information services are of major importance since they bring order into the chaos of a distributed, heterogeneous information space. In the regional healthcare network of Crete, a wide range of regional services, which operate as enabling middleware services [47], are being developed. WebOnCOLL will exploit these services as they become available.

1) Resource Services: Regional resource services provide availability information on physical resources, such as hospital departments, diagnostic modalities, mobile emergency units, and their characteristics. In this way, regional resource services

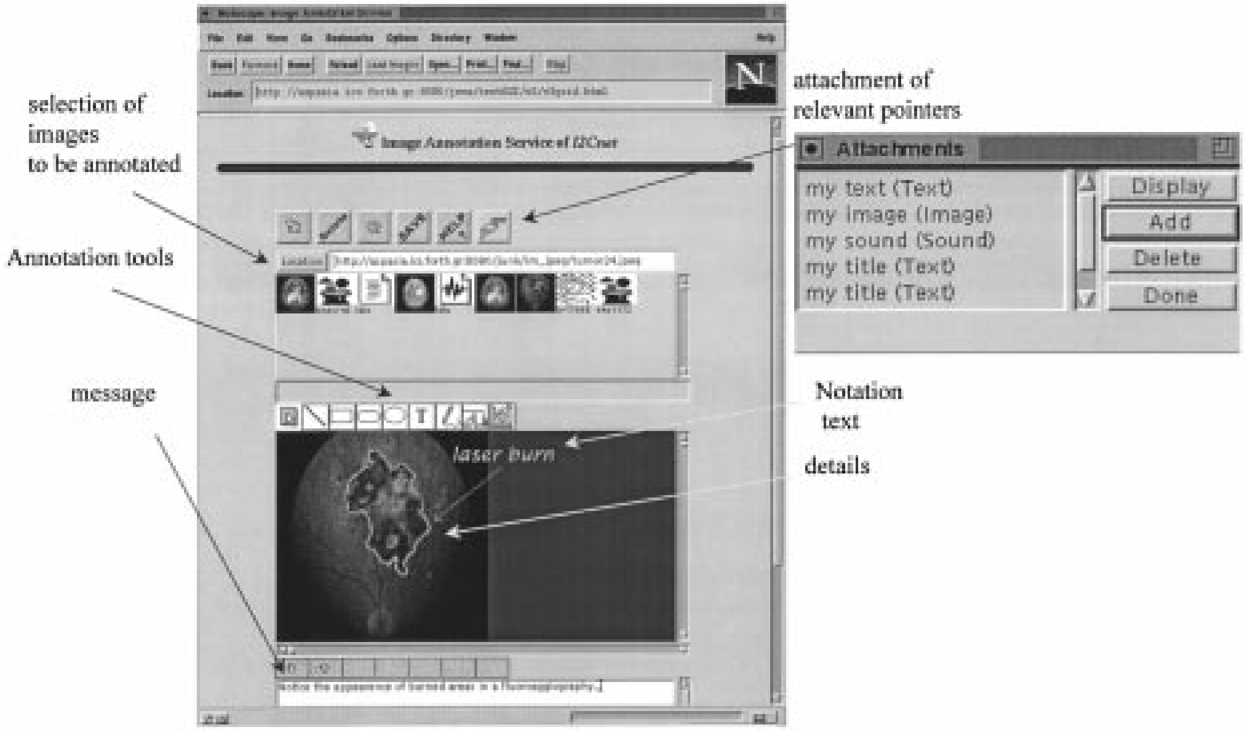

Fig. 5. The annotation service facilitates the interaction of healthcare professionals with diagnostic images and patient record data.

will facilitate access and booking of appointments with the appropriate healthcare facilities. Educational resources, such as public health information and access to digital medical libraries will also be available. In addition, links to resource services of other regions strengthen transregional collaboration.

2) Directory Services: Directory services maintain indexed information on healthcare record segments and yellow pages for healthcare facilities and human resources. For healthcare professionals, directory services typically provide information on contacts, roles, and authorities. In addition, the storage of public keys in the directory enables confidential and authenticated communication.

3) Security and Confidentiality Services: Security and confidentiality services are based on a regional certification authority, which will provide digital certificates to healthcare facilities and human resources. The purpose of the certification authority is to certify the role and authority of the collaborators. The combination of digital signatures for authentication, public key cryptography for recipient authentication and Secure Socket Layer (SSL) for secure data-transfer, provide the technological framework for secure collaboration. However, the policies and procedures related to granting authorization to a healthcare actor executing specific medical acts needs to be addressed on a per case basis. In this way, a trust infrastructure [54] will evolve in the regional healthcare network of Crete, which will interoperate with the collaboration infrastructure of WebOnCOLL.

IV. Architecture of WebOnCOLL Servers

A basic requirement of the WebOnCOLL architecture is conformance to the so-called "Open Workspace" paradigm [36], [26]. According to the Open Workspace paradigm, the computer should be integrated into the workspace, rather than isolating the user from the workspace. To attain this goal, the WebOnCOLL architecture is centered on the concept of virtual workspaces and user profiles. Virtual workspaces provide the collaboration framework, while user profiles customize it to the needs and tasks of the current user.

The architecture of a WebOnCOLL server appears in Fig. 6. The basic components of the architecture are the web server, the workspace manager, the user profile manager, and the information repositories, which reside on a database management system (DBMS) and the file system. The workspace manager controls virtual workspaces and provides notification and awareness information to its clients. The user profile manager maintains user profiles and personalized information channels and customizes the collaboration environment to user preferences.

A. Workspace Manager

The workspace manager is a Java application, which runs on the WebOnCOLL server. It communicates using TCP/IP with client workspace applets, which run on the web browser of the user. The workspace applet retrieves the contents of the workspace from the web by using the HTTP protocol to take advantage of the client-configured proxy cache hierarchy.

For each workspace, the workspace manager maintains access control information, an event log, and a list of the currently connected users. Each workspace corresponds to a directory on the file system of the WebOnCOLL server. This workspace directory contains a file called "contents," which specifies the data collection associated with the virtual

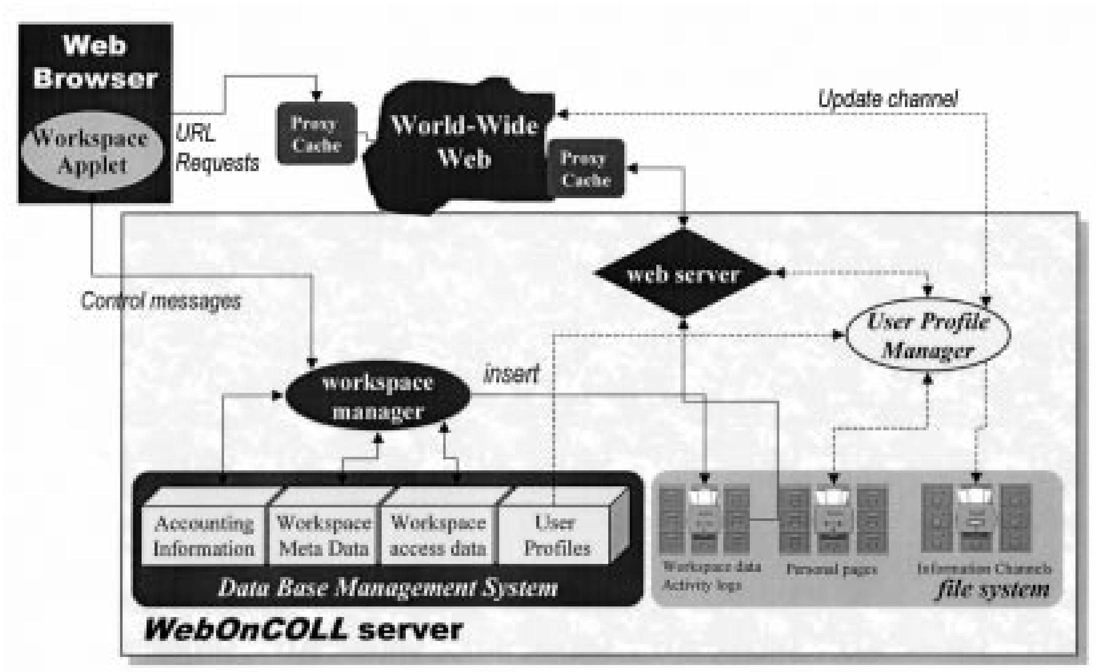

Fig. 6. Architecture of a WebOnCOLL server.

workspace. For each object in the collection, the "contents" file specifies *name*, *URL*, *type*, *actor*, and *date*. Name is the display name of the object, type is the media type it contains, actor is the user that inserted the object in the collection, and date is the date when it was inserted. The workspace directory also contains all local objects of the workspace, which may be products of service requests, such as query responses or annotation objects. For local objects, the workspace maintains additional property objects, which contain information specific to the object. For example, if a workspace object is an annotation, the property object maintains meta-information about the annotation, i.e., author, date, and keywords.

As shown in Fig. 7, the communication of the workspace client applet with the workspace manager involves administrative, awareness, and notification messages. In response to administrative messages, the workspace manager updates the workspace. In response to notification and awareness messages, the workspace manager provides contents, notification, or awareness information, as appropriate. The client applet has two threads of control, the *user thread* that responds to user requests and passes administrative requests to the workspace manager and the *update thread*, which polls the workspace manager for updates in the status of the workspace. The update thread runs in the background, with low priority, so that the user does not experience any response delays.

1) Workspace Applet Initialization: The workspace manager listens on a preset TCP/IP port and creates a new thread for each request of a client workspace applet. In its initialization phase, a workspace applet opens a connection to the workspace manager and provides authentication information, i.e., *workspace*, *user*, and *password*. The workspace manager verifies that the user has indeed the authority to access the workspace and records the connection in the *event log*. Then, the workspace manager adds the user to the list of users that are currently connected to the requested workspace. Finally, the workspace manager returns to the client applet a *session identifier* (sessionID), the workspace contents, and the date the workspace was last modified. The applet receives the workspace contents and creates an internal list for the collection. Then, for each object in the collection, it displays its name and an icon suitable for the media type of the object (see Fig. 3). Only when the user requests to view a specific object is the object fetched over the web by using the HTTP protocol. Note that alternative interfaces and presentations of the workspace contents are possible. For example, we have considered an XML [5] interface to a virtual workspace.

When the client workspace applet exits, it sends a *disconnect* message to the workspace manager, which removes the user from the list of users currently connected to the workspace.

2) Administrative Messages: The workspace applet may send to the workspace manager a number of administrative messages: *insert_object*, *delete_object*, *create_object*, *service*, and *post*. Each message consists of a header specifying the attributes of the requester and a body containing the data for the request. The workspace manager parses the message and services the request. When the service execution is complete, the workspace manager sends an acknowledgment followed by the result of the request or an error code.

The *insert_object* message has four attributes: sessionID, name, URL, and type. In response to the *insert_object* message, the workspace manager checks the user permissions, updates the event log, updates the workspace, and returns the revised workspace contents or an error code to the client applet.

The *delete_object* message is analogous to the *insert_object*. The applet sends to the workspace manager the sessionID and the name of the workspace object to be deleted. Again

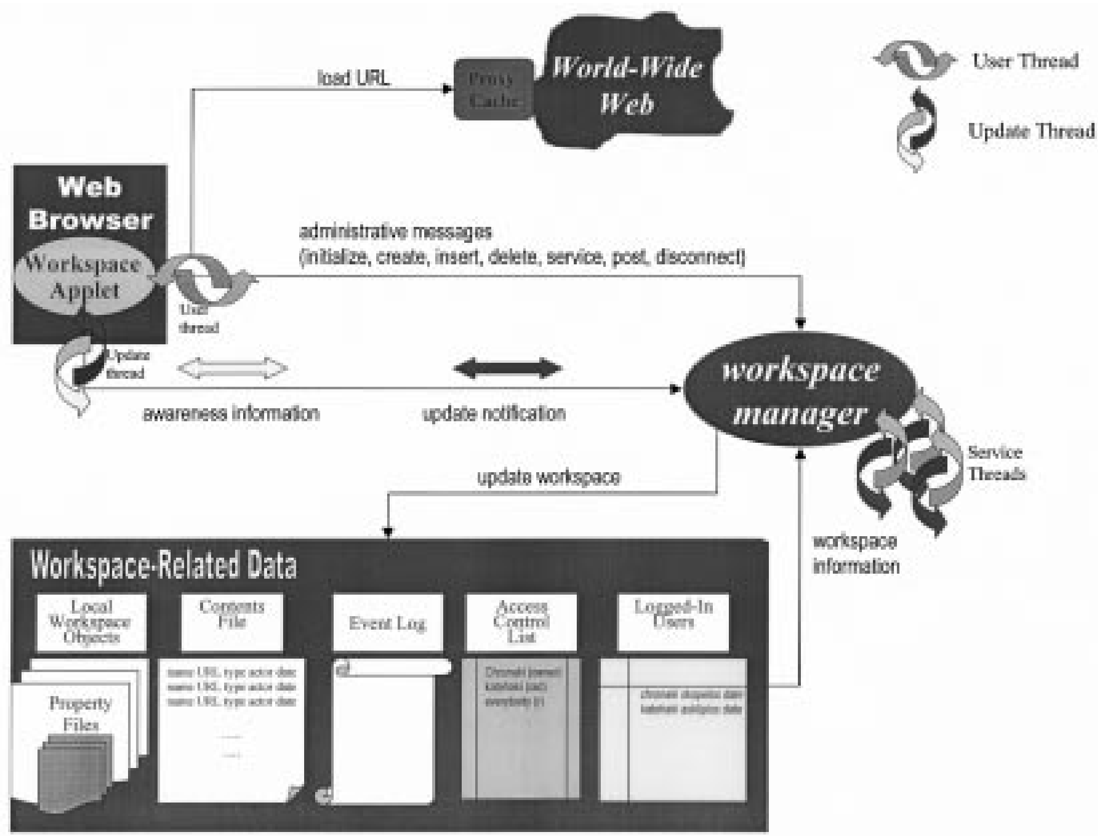

Fig. 7. The workspace manager is responsible for maintaining workspace-related data. The client workspace applet has two threads of control. The user thread of control communicates administrative information to the workspace manager, while the update thread is responsible for update notifications and user awareness.

the workspace manager checks the permissions of the user, updates the event log, updates the workspace, and returns the revised workspace contents or an error code to the client applet.

The *create_object* message applies to the case where the client workspace code is part of a service applet, as in the case of the annotation service. The annotation applet creates an annotation in memory, and the user has the option of saving the annotation in the current workspace. The applet sends to the workspace, the sessionID, the name of the file to be created, and the contents of the file. Upon receipt, the workspace manager checks the permissions of the user, updates the event log, creates a new file with the provided contents in the workspace directory, and updates the workspace. Then, it returns the revised workspace contents to the client applet.

The *service* and *post* messages refer to requesting the execution of a service (job) on the WebOnCOLL server and posting an object to a bulletin board, respectively. Handling a service request is described in the section on session management (Section IV-A4). In case of a post request, the workspace manager checks that the user has the right to post to the target bulletin board and updates the corresponding workspace.

3) Notification and Awareness: There are two types of changes that may occur in the state of the workspace: 1) the contents of the workspace have changed and 2) the list of the users that are connected to the workspace have changed. The update thread of a client workspace applet periodically queries the workspace manager for updates in the state of the workspace, sending a *notify_me* message to the browser with attribute sessionID. The workspace manager updates the event log and checks if the state of the workspace has been modified since the specific client was last notified. If the state of the workspace has been modified, the workspace manager responds with the new state of the workspace.

Note that this notification and awareness policy enforces an information pull, rather than information push approach. This design choice was based on the simplicity of the implementation. We plan to experiment with distributed objects (i.e., CORBA [51], [52]) in the communication of the client workspace applet and the workspace manager. In that context, the event and notification services of CORBA will ease the implementation of information push.

4) Session Management: A client workspace applet may request the execution of a service (job) on the WebOnCOLL server. Such a job may involve the update of an information channel, an Internet search, or a batch copy of a large number of files. The duration of such jobs may be longer than the duration of a specific session. Therefore, the workspace manager mediates the execution of such jobs. As shown in Fig. 8, the workspace manager maintains a list of active users. For each active user, a data structure maintains jobs in progress. For each service request, the workspace manager spawns a new thread that monitors the job and updates the data structure. When the job ends, account information is stored in the event log and the results of the job are stored in the corresponding workspace.

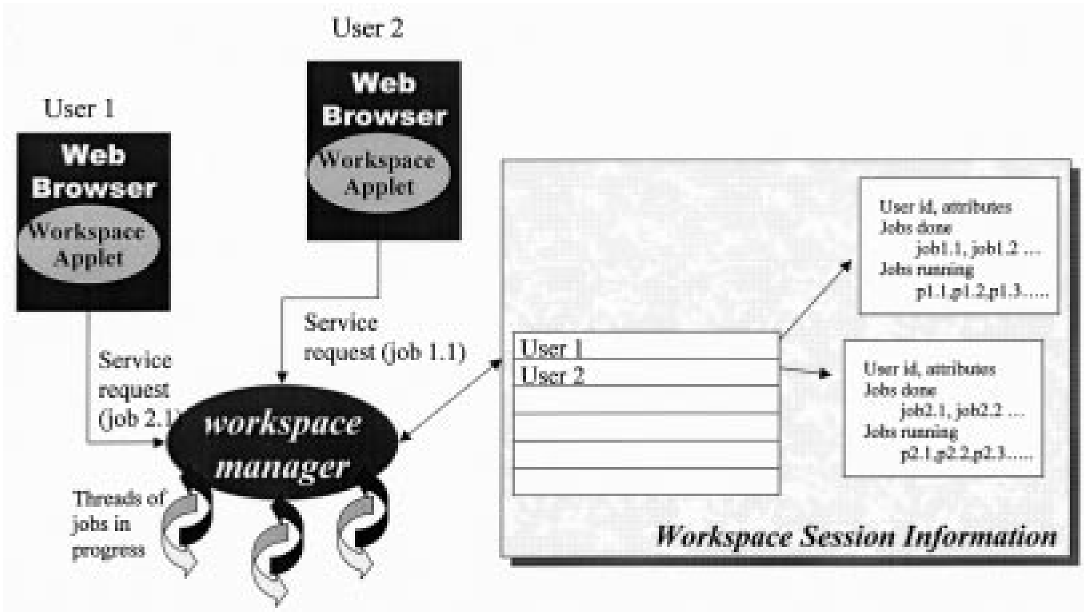

Fig. 8. Session information in the workspace manager.

B. User Profile Manager

An environment for collaboration should not impose a particular mode of interaction. It should, instead, provide a wide range of capabilities that can be quickly and easily selected and configured to the task at hand. One way to attain this goal is to adapt the behavior of the system to the nature of the collaboration and user preferences. In this context, interface agents are employed to customize the application interface to the evolving needs and preferences of users. Typically, interface agents employ machine learning technology to learn from user behavior and create the user profile [45], [46].

Currently, user profiles are increasingly used to customize an application environment to the needs of individual users. Fourth generation browsers maintain advanced user profiles, which dramatically change the behavior of the application. These user profiles include e-mail profiles, address books, bookmarks, news profiles, Internet history, cache proxy configuration, directory services, certificates, channels, helper applications, and user interface preferences. User interface preferences include fonts, colors, style sheets, start-up, and language configuration.

A user profile in WebOnCOLL is an extended version of a user profile in a fourth generation browser. Additional information maintained by user profiles in WebOnCOLL concerns account information and subscribed services. Account information includes information on the preferred application environment (user preferences, virtual workspaces, etc.), mobility information, such as current Internet service provider, and sensitive information, such as access codes. Subscribed services relate to personalized information services, such as notification services and information channels.

The user profile manager currently being implemented employs user profiles to customize the appearance of personal pages and provide access to data collected off-line, notifications, and alerts. The maintenance of information channels by the user profile manager involves interaction with task agents that collect information using push and "smart-pull" technology [48].

1) Information Channels: In the context of information channels, WebOnCOLL downloads up-to-date, personalized information in the background and stores it in off-line caches, which are accessed by the user as specialized virtual workspaces. This allows users to experience dynamic content without waiting for pages to load by using a technology that enables the push delivery of information without requiring a user to manually request or search for it. In WebOnCOLL, an information channel may be associated with personal web pages, bulletin boards, discussion lists, shared workspaces, and medical case folders. Furthermore, a number of issues related to information filtering remain to be investigated [55], [56].

V. FUTURE WORK

Our future work in WebOnCOLL will focus on two separate web-based collaboration projects: 1) developing a collaboratory for health monitoring in primary care and 2) implementing a virtual medical office for teleworking.

The primary objective of the collaboratory for health monitoring in primary care is to set up the framework for the continuous monitoring and evaluation of health data and indicators at a regional level. Health monitoring in a regional healthcare network is a complex, collaborative activity, which requires the synergy of scientific, administrative, and technical actors. In this project, healthcare researchers from different European regions, review and analyze health indicators, based on primary care data, to determine their relevance, quality, and coverage. The data collection process will involve the primary healthcare centers in the region of Crete. Appropriate information channels linked to discussion lists, personal pages, and bulletin boards will be set up to facilitate the exchange of ideas regarding methods, rules, and procedures for the routine collection, analysis, and reporting of primary data. Population health indicators as well as calendars, meetings, reports, and recommendations will be available on the web site. Directory services will provide access to the participating research, administrative, and technical staff in a secure and reliable way.

The second project applies to medical professionals who need constant access to their working environment, even though this environment is distributed [24], [42], [50]. Consider a medical professional who maintains a university office, a home office, and a private practice office. There is constant interaction with assistants, colleagues, and patients in all offices, and the professional needs to access information (colleagues, medical cases, reminders, notifications, etc.) regardless of location. The WebOnCOLL collaboration environment will be used to create a virtual medical office, which will improve the ability of medical professionals to work while enhancing their mobility.

VI. Usability and Acceptance Issues

To create a usable environment for collaboration, you need content, information processing methods, and technology. In the area of content, healthcare professionals need access to a wide variety of information, including medical literature, expert summaries as found in textbooks and guidelines, and information on medications and diagnostic tests [35]. Information processing methods are necessary to find the relevant content. Finally, technology is necessary to deliver current, well-structured, high-quality content in a timely manner. High-speed data networks, standard protocols, an open systems architecture, and cross-platform applications are necessary technological components.

The major problem with a collaborative environment, as with any other form of technology, is that no matter how advanced a system may be, it is the users that will have the final word, accepting or discarding it. Grudin [34] cites various reasons for the failure of CSCW systems; one of the principal ones being failure in obtaining a critical mass of users. If there is a small number of users, the cost for each user who uses the system is likely to exceed the benefit. Only when there is a sufficient number of users does the benefit exceed the cost. When the benefit starts to become visible, users will start using it and new ones will join.

Like clinical information systems, collaboration services alter traditional practice patterns [30]. Clinical information systems impose major limitations on how clinical data are recorded and how the medical record is organized. In the same way, collaboration services, in the wide sense, interfere with the way that physicians organize their thought process in caring for patients. Furthermore, collaboration environments like WebOnCOLL will affect professional relationships between individuals and groups within healthcare organizations. The response of physicians is shaped by their perception of how these new practices will affect these relations and the accomplishment of professional goals.

Methods that have been used in attempts to affect physicians practice behavior [30] include 1) the provision of information about their performance compared with norms or with other physicians, 2) the direct involvement of physicians in the effort by the use of opinion leaders among them in translating medical advances into practice, and 3) the identification of specific benefits to individuals and organizational groups. The only viable way for the "real-world" deployment of collaboration services is by paying attention to social and organizational factors. Given that these collaboration services affect professional activities, their impact on practice patterns and professional relationships within healthcare organizations should be studied.

VII. Conclusion

Healthcare is an important application domain in the emerging information society. In recent years, we have all witnessed the gradual transformation of health informatics into health telematics, a process that still continues. For this transformation to be successful, a strategy is needed for the creation of an integrated healthcare information infrastructure. Furthermore, in developing such a strategy, the utilization of this infrastructure to *provide clinically significant added-value services that would justify its cost and ultimately benefit the patient* must be considered [47].

Computer-supported collaboration enhances communication between medical personnel and their patients. On the other hand, information access that delivers information in more complex representations, like sounds and images, requires enormous amounts of data storage and network capacity. Thus, in order for the new environment to be accepted, it must improve an individual's ability to work. Hence, advanced network infrastructures are vital for supporting network connection and fast access times. Cross-platform whiteboards, shared workspaces, and application sharing that utilize less bandwidth than that required by video have been implemented, but remain at an early stage of development, especially with regard to interoperability. Several more years of research will be needed until the collaboratory tools are mature enough to be acceptable for end usage.

The web-based medical collaboration services of WebOnCOLL support asynchronous user collaboration through the use of virtual workspaces, which implement collaboration concepts, such as personal web pages, bulletin boards, discussion lists, shared workspaces, and medical case folders. The annotation service allows users to interact with the content of information repositories and create annotations that may be e-mailed to any user of the regional network, be posted to various fora, or saved in a private collection or shared workspace. The on-line collaboration service permits user interaction in pseudoreal time through the discussion of multimedia data objects, ranging from ECG's and diagnostic images to laboratory results.

The deployment of the WebOnCOLL environment in the context of several projects relevant to health monitoring, teleworking, and education aims to promote various forms of medical collaboration. In this way, not only the interaction of a large number of users with the available information repositories will be facilitated, but also the active participation of users in the evolution of the regional healthcare network will be promoted.

Acknowledgment

The authors would like to thank G. Halkiadakis, S. Kostomanolakis, V. Spitadakis, and N. Stathiakis for constructive comments made in various stages of this work.

References

[1] A. Dix, "Challenges and perspectives for cooperative work on the web," in *Proc. ERCIM Workshop CSCW Web*, Sankt Augustin, Germany, Feb. 7–9, 1996, http://www. hud.ac.uk/schools/comp+maths/private/alan/papers/CSCW+WWW/issues.html.

[2] G. Calvary, J. Coutaz, and L. Nigay, "From single-user architectural design to PAC*: A generic software architecture model for CSCW," in *Proc. CHI'97*, htpp://ww.acm.org/sigchi/chi97/proceedings/paper/jcc.htm.

[3] C. Ellis and J. Wainer, "A conceptual model of groupware," in *Proc. CSCW'94, ACM Conf. Comput.-Supported Cooperative Work*, R. Furuta and C. Neuwirth, Eds. 1994, pp. 79–88, http://www.acm.org/pubs/articles/proceedings/cscw/192844/p79-ellis/p79-ellis.pdf.

[4] W3C, "HTML 4.0 specification," Proposed Recommendation, Nov. 7, 1997, http://www.w3.org/TR/PR-html40/.

[5] D. Connolly, "XML: Principles, tools, and techniques," *WWW J.*, vol. 2, Fall 1997, http://www.ora.com/catalog/ wjfall97/index.html.

[6] R. Bentley, U. Busbach, and K. Sikkel, "The architecture of the BSCW shared workspace system," in *Proc. ERCIM Workshop CSCW Web*, Sankt Augustin, Germany, Feb. 7–9, 1996, http://bscw.gmd.de/Papers/.

[7] K. Holtman, "The FUTPLEX system," in *Proc. ERCIM Workshop CSCW Web*, Sankt Augustin, Germany, Feb. 7–9, 1996, http://gewis.win.tue.nl/∼ koen/futplex_CSCW.html.

[8] P. De Bra, "DreSS: A simple document repository service station," in *Proc. ERCIM Workshop CSCW Web*, Sankt Augustin, Germany, Feb. 7–9, 1996, http://orgwis.gmd.de/projects/W4G/presentations/dress.html.

[9] R. T. Kouzes, J. D. Myers, and W. A. Wulf, "Collaboratories: Doing science on the Internet," *IEEE Comput.*, vol. 29, pp. 40–46, Aug. 1996, http://www.wvu.edu/∼research/DOE/IEEEcollaboratory.html.

[10] J. Grudin, "Computer-supported cooperative work: History and focus," *IEEE Comput.*, vol. 27, pp. 19–26, May 1994.

[11] J. Rhyne and C. Wolf, "Tools for supporting the collaborative process," in *Proc. UIST'92, 5th Ann. ACM Symp. User Interface Softw. Technol.*, pp. 161–170.

[12] ITU, "X.500 recommendations," http://www.itu.int//itudoc/itu-t/rec/x.html.

[13] W. Yeong, T. Howes, and S. Kille, "RFC 1777: Lightweight directory access protocol," Mar. 1995, http://ds.internic.net/rfc/rfc1777.txt.

[14] H.320 standards for video conferencing.

[15] W. Gramlich, "Public annotations," *Workshop WWW Collaborative Work*, Cambridge, MA, Sept. 11–12, 1995, http://www.w3.org/ Collaboration/Workshop/Proceedings/P10.html.

[16] J. Clyman *et al.*, "IE4 vs. communicator: Face off," *PC Mag.*, Special Issue on "The Future of the Web," vol. 16, no. 20, 1997, http://www.zdnet.com/pcmag/features/browsers40/_open.htm.

[17] V. Jagannathan *et al.*, "An overview of the CERC Artemis project," Concurrent Eng. Res. Center, West Virginia Univ., Morgantown, Tech. Rep. CERC-TR-RN-95-002, Apr. 1995, http://www.cerc.wvu.edu/nlm/artemis/summary.html.

[18] Distributed Collaboratory Experiment Environments, http://www.itg.lbl.gov/∼jtchew/DCEE_overview.html.

[19] Collaboratory for Research on Electronic Work, Univ. Michigan, Ann Arbor, http://www.crew.umich.edu/CREW.home.html.

[20] S. E. McDaniel, "Providing awareness information to support transitions in remote computer-mediated collaboration," presented at *Doctoral Consort., CHI'96*, http://www.acm.org/sigchi/chi96/proceedings/doctoral/McDaniel_Susan /sem_doc.html.

[21] InterMed Collaboratory, http://smi-web.stanford.edu/projects/intermed-web/Overview.html.

[22] D. G. Kilman and D. W. Forslund, "An international collaboratory based on virtual patient records," *Commun. ACM*, vol. 40, pp. 111–117, Aug. 1997.

[23] Virtual telemedicine Office, http://www.telemedical.com.

[24] G. Marcelino, "The virtual and the paperless office," *IEEE Comput.*, vol. 28, pp. 120, June 1995.

[25] K. Kamiya, M. R scheisen, and T. Winograd, "Grassroots: A system providing a uniform framework for communicating, structuring, sharing information, and organizing people," in *Proc. 5th Int. WWW Conf.*, Paris, France, May 6–10, 1996, http://www.acm.org/sigchi/chi96/proceedings/shortpap/Kamiya/kk_txt.html.

[26] M. Billinghurst, S. Baldis, E. Miller, and S. Weghorst, "Shared space: Collaborative information spaces," in *Proc. CHI'97*, http://www.hitl.washington.edu/publications/p-96-5.

[27] K. Andrews, F. Kappe, and H. Maurer, "Serving information to the web with Hyper-G," in *3rd Int. WWW Conf.*, Darmstadt, Germany, 1995, http://www.igd.fhg.de/www/www5/papers/105/hgw3.html.

[28] E. Leisch, S. Sartzetakis, M. Tsiknakis, and S. C. Orphanoudakis, "A framework for the integration of distributed autonomous healthcare information systems," *Med. Inform.*, to be published, http://www.ics.forth.gr/ICS/acti/cmi_hta/publications/mi97int/mi97int.html.

[29] M. Blazantonakis, V. Moustakis, and G. Charissis, "Seamless care in the health region of Crete: The STAR* case study," in *Proc. MIE'97*, Porto Carras, Greece, May 1997, http://www.ics.forth.gr/ICS /acti/cmi_hta/publications/papers.html.

[30] J. G. Anderson, "Clearing the way for physicians' use of clinical information systems," in *Commun. ACM*, vol. 40, pp. 83–90, Aug. 1997.

[31] C. E. Chronaki, X. Zabulis, and S. C. Orphanoudakis, "I^2Cnet medical image annotation service," *Med. Inform.*, to be published, http://www.ics.forth.gr/ICS/acti/cmi_hta/publications/eup96ann /eup96ann.html.

[32] S. C. Orphanoudakis, C. E. Chronaki, and D. Vamvaka, "I^2Cnet: Content-based similarity search in geographically distributed repositories of medical images," *Comput. Med. Imaging Graph.*, vol. 20, no. 4, pp. 193–207, 1996, http://www.ics.forth.gr/ICS/acti/cmi_hta/publications/dglib97/dglib97.html.

[33] M. Zikos, E. Kaldoudi, and S. C. Orphanoudakis, "DIPE: A distributed environment for medical image processing," in *Proc. MIE'97*, May 25–29, 1997, pp. 465–469, http://www.ics.forth.gr/ICS/acti/cmi_hta/publications/eup96comed/eup96comed.html.

[34] J. Grudin, "Why CSCW applications fail: Problems in the design and evaluation of organizational interfaces," in *CSCW'88 Proc. Conf. Comput.-Supported Cooperative Work*, ACM SIGCHI & SIGOIS, pp. 85–89.

[35] W. M. Detmer and E. H. Shortliffe, "Using the internet to improve knowledge diffusion in medicine," *Commun. ACM*, vol. 40, pp. 101–108, Aug. 1997.

[36] H. Ishii and N. Miyake, "Toward an open shared workspace: Computer and video fusion approach of team workstation," *Commun. ACM*, vol. 36, pp. 53–62, July 1993.

[37] T. Dorcey, "CU-seeme desktop videoconferencing software," *Connexions*, vol. 9, Mar. 1995; http://cu-seeme.cornell.edu/DorceyConnexions.html.

[38] W. Reinhard, J. Schweitzer, G. Vlksen, and M. Weber, "CSCW tools: Concepts and architectures," *IEEE Comput. Mag.*, vol. 27, pp. 28–36, May 1994.

[39] O. Sandor and K. Tollman, "@Work: The design of a new communication tool," in *Proc. ERCIM Worksjop CSCW Web*, Sankt Augustin, Germany, Feb. 7–9, 1996.

[40] M. Roscheisen, T. Mogensen, and T. Winograd, "Beyond browsing: Shared comments, SOAP's, trails, and on-line communities," in *Proc. 3rd Int. WWW Conf.*, Darmstadt, Germany, Apr. 1995, http://www-diglib.stanford.edu/diglib/pub/reports/ bro_www95.html.

[41] M. Roscheisen, C. Mogensen, and T. Winograd, "Shared web annotations as a platform for third-party value-added information providers: Architecture, protocols, and usage examples," Stanford Integrated Digital Library Project, Comput. Sci. Dept., Stanford Univ., Stanford, CA, Tech. Rep. STAN-CS-TR-97-1582, Apr. 1995, http://www-diglib.stanford.edu/diglib/pub/reports/commentor.html.

[42] M. Sohlenkamp and G. Chwelos, "Integrating communication, cooperation, and awareness: The DIVA virtual office environment," in *Proc. CSCW'94*, pp. 331–343, http://www.acm.org/pubs/articles/proceedings /cscw/192844/p331-sohlenkamp/p331-sohlenkamp.pdf.

[43] S. Whittaker and S. Candace, "Information overload leads to overloaded e-mail: An empirical study of e-mail usage," in *Proc. CHI-96*, pp. 276–283, http://www.lotus. com/lotus/research/.

[44] M. Bainbridge, P. Salmon, A. Rappaport, G. Hayes, J. Williams, and S. Teasdale, "The problem oriented medical record—Just a little more structure to help the world go round," http://ww.ncl.ac.uk/∼nphcare /PHCSG/conference/camb96/mikey.html.

[45] P. Maes, "Agents that reduce work and information overload," *Commun. ACM*, vol. 37, pp. 31–40, July 1994; http://pattie.www.media.mit.edu/people/pattie/CACM-94/CACM-94.p1.html.

[46] Y. Lashkari, M. Metral, and P. Maes, "Collaborative interface agents," in *Proc. 12th Nat. Conf. Artificial Intell.*, vol. 1. Seattle, WA: AAAI, Aug. 1994; http://agents. www.media.mit.edu/groups/agents/papers/aaai-ymp/aaai.html.

[47] M. Tsiknakis, C. E. Chronaki, S. Kapidakis, C. Nikolaou, and S. C. Orphanoudakis, "An integrated architecture for the provision of health telematic services based on digital library technologies," *Int. J. Digital Libraries*, Special Issue on, "Digital Libraries in Medicine," vol. 1, no. 3, pp. 257–277, 1997; http://www.ics.forth.gr/ICS/acti/cmi_hta/publications/dglib97/dglib97.html.

[48] W3C, in *Proc. W3C Workshop Push Technol.*, Boston, MA, Sept. 8–9, 1997; http://www.w3.org/Architecture/9709_Workshop/.

[49] M. A. Schickler, M. S. Mazer, and C. Brooks, "Pan-browser support for annotations and other meta-information on the world wide web," in *Proc. 5th Int. WWW Conf.*, Paris, France, May 6–10, 1996.

[50] N. Streitz, J. Geiler, J. Haake, and J. Hol, "DOLPHIN: Integrated meeting support across local and remote desktop environments and liveboards," in *Proc. CSCW'94*, pp. 345–358; http://www.acm.org/pubs/citations/proceedings/cscw/192844/p345-streitz/.

[51] D. C. Schmidt and S. Vinoski, "The OMG events service," *Object Interconnections*, column 9, SIGS C++ Rep., Feb. 1997.

[52] J. Siegel, "CORBA: Fundamentals and programming," in *OMG*. New York: Wiley, 1996.

[53] M. Zikos, S. Stephanidis, and S. C. Orphanoudakis, "CoMed: Cooperation in medicine," in *Proc. EuroPACS'96*. Heraklion, Crete, Greece, pp. 88–92, http://www.ics.forth.gr/ICS/acti/cmi_hta/publications/eup96comed/eup96comed.html.

[54] Y.-H. Chu, J. Feigenbaum, B. LaMacchia, P. Resnick, and M. Strauss, "REFEREE: Trust management for web applications," in *Proc. 6th Int. WWW Conf.*, 1997, http://www. w3.org/PICS/TrustMgt/doc/referee-WWW6.html.

[55] M. Q. W. Baldonado and T. Winograd, "SenseMaker: An information-exploration interface supporting the contextual evolution of a user's interests," in *Proc. CHI'97, Conf. Human Factors Comput. Syst. Hyatt Regency Atlanta Hotel*, Atlanta, GA, http://www.acm.org/sigchi/chi97/proceedings/paper/mwb.htm.

[56] R. Barrett, P. P. Maglio, and D. C. Kellem, "How to personalize the web," in *Proc. CHI'97, Conf. Human Factors Comput. Syst. Hyatt Regency Atlanta Hotel*, Atlanta, GA, http://www.acm.org/sigchi/chi97/proceedings/paper/rcb-wbi.htm.

[57] E. Kaldoudi, M. Zikos, E. Leisch, and S. C. Orphanoudakis, "Agent-based workflow processing for functional integration and process re-engineering in the health care domain," in *Proc. EuroPACS'97*, Pisa, Italy, pp. 247–250, http://www.ics.forth.gr/ICS/acti/cmi_hta/publications/eup97wf/eup97wf.html.

Catherine E. Chronaki received a degree in computer engineering from the University of Patras, Patras, Greece, and the M.S. degree in computer science from the Rochester Institute of Technology and the University of Rochester, both located in Rochester, NY.

She has been with the Center for Medical Informatics and Health Telematic Application, Institute of Computer Science, Foundation for Research and Technology (FORTH), Crete, Greece, since 1992. Her research interests include, but are not limited to, medical information systems, Internet applications, integration architectures, and digital medical libraries.

Dimitrios G. Katehakis (S'88–M'92) received a degree in electrical engineering from the Technical University of Patras, Patras, Greece, in 1991 and the M.S. degree in electrical engineering from the University of Maryland, College Park, in 1993.

He is a Telecommunications Engineer at the Center for Medical Informatics and Health Telematics Applications, Institute of Computer Science, Foundation for Research and Technology (FORTH), Crete, Greece. He works on data integration analysis and the design of object-oriented systems and distributed architectures for medical telematics applications across interregional networks. Currently, his research focuses on web-based medical collaboration environments as well as the intelligent management and communication of multimedia data.

Xenophon C. Zabulis received a degree in computer science from the University of Crete, Crete, Greece, in 1991.

He is a Postgraduate Student at the Center for Medical Informatics and Health Telematics Applications, Institute of Computer Science, Foundation for Research and Technology (FORTH), Crete. His current research interests are in the areas of intelligent agents and distributed systems.

Manolis Tsiknakis (M'89) received the B.Eng. degree in electronic engineering, the M.Sc. degree in microprocessor engineering, and the Ph.D. degree in control systems engineering from the University of Bradford, Bradford, U.K.

He was a member of the Research Staff at the Department of Control Systems Engineering, University of Bradford, from 1985 to 1990. In 1992, he joined the Institute of Computer Science, Foundation for Research and Technology (FORTH), Crete, Greece, 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.

Stelios C. Orphanoudakis (S'72–M'76–SM'83) received the B.A. degree (magna cum laude with highest distinction) in engineering sciences from Dartmouth College, Hanover, NH, the M.S. degree in electrical engineering from the Massachusetts Institute of Technology, Cambridge, and the Ph.D. degree in electrical engineering from the Thayer School of Engineering, Dartmouth College.

He was with the Departments of Diagnostic Radiology and Electrical Engineering, Yale University, New Haven, CT, from 1975 to 1991. He is currently Director of the Institute of Computer Science, Foundation for Research and Technology (FORTH) and Professor of Computer Science, University of Crete, Greece. He has many years of academic and research experience in the fields of computer vision and robotics, intelligent image management and retrieval by content, medical informatics, and medical imaging.

Dr. Orphanoudakis is a member of many honorary and professional societies. 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 Board of Directors of the EuroPACS society and of the National Telecommunications Commission and the National Advisory Research Council of Greece.