

The I^2Cnet Service Architecture Paradigm

Xenophon Zabulis^{1,2}, Catherine Chronaki²,
and Stelios Orphanoudakis^{1,2}

¹ *Department of Computer Science, University of Crete*
² *Institute of Computer Science, FO.R.T.H., Heraklion, Crete, Greece*

pp. 596-600

MIE'97

Medical Informatics Europe
14th International Congress, Porto Carras, Greece

May 25-29, 1997

The I^2Cnet Service Architecture Paradigm

Xenophon Zabulis^{1,2}, Catherine Chronaki², and Stelios Orphanoudakis^{1,2}

¹*Department of Computer Science, University of Crete*

²*Institute of Computer Science, FO.R.T.H., Heraklion, Crete, Greece*

Abstract. The main objective of the Image Indexing by Content network (I^2Cnet) is to provide network-transparent content-based access to medical image archives based on a collection of interoperable Internet/intranet added-value services. This paper discusses I^2Cnet , focusing on its service architecture paradigm. I^2Cnet services such as image annotation, processing, description, and content-based retrieval, as well as the on-line collaboration service are presented. Exemplary user sessions are used to illustrate how virtual workspaces facilitate the interoperation of I^2Cnet services, following the “network computer” approach to information management.

1. Introduction

I^2Cnet [1] (Image Indexing by Content network) is an on-going project at ICS-FORTH, whose goal is to provide network-transparent content-based access to medical image archives as an Internet/intranet added-value service. Through a standard Web browser, healthcare professionals will be able to interact with image collections, browse images similar to a query image, compare these images to images from other collections, and contribute their own images or comments. Specific I^2Cnet services currently available on the world-wide Web include image processing, content- and annotation-based search for images and image-related information, and authoring of annotations and image descriptions (visit: <http://www.ics.forth.gr/~telemed/services.html>). I^2Cnet is organized as a network of I^2C servers and brokers. Each I^2C server distributes its own set of services to the brokers, which are responsible for updating the I^2Cnet service directory. I^2C brokers maintain directory information and service profiles based on user-feedback and access patterns and use this information to handle network-transparent service requests. In this environment, virtual workspaces maintain the common context in the form of a heterogeneous data collection which points to all the relevant data objects. Virtual workspaces persist between user sessions, among services (service interoperability), and among users (user cooperation).

The rest of the paper is organized as follows: Section 2 elaborates on the service architecture paradigm followed by I^2Cnet . Section 3 describes the services currently supported by I^2Cnet and section 4 illustrates exemplary user sessions in I^2Cnet . Finally, section 5 concludes the paper.

2. I^2Cnet Service Architecture Paradigm

The term service architecture originated from the communications field, in which telephone and cable companies offer services of various levels of quality to their customers. The same concept fits very well with the requirements defined for the I^2Cnet architecture:

1. all services are accessible to authorized users anywhere in the world from any Java-enabled browser running on a computer with Internet access.
2. various users attain different privileges towards information access and service quality
3. each I^2C server maintains an autonomous repository of images and related data and offers a potentially different set of services
4. users may specify the server or servers which should handle a service request, or may request a particular service by name (network-transparency)
5. network-transparent services take user feedback and user access patterns into account
6. I^2Cnet services interoperate with the regional healthcare network.

These complex requirements call for a network-centric architectural design which is loosely-coupled. Traditional client-server architectures are more closely-coupled in the sense that the client is affected by changes in the server. Furthermore, to be able to use an application on a specific platform, the

installation of the appropriate client software is necessary. In such a setting, *accessibility*, *portability*, and *software evolution* entail considerable effort.

In contrast, the service architecture of I^2Cnet deals comfortably with these issues. *Accessibility* is limited only by access rights, and the availability of an Internet connection. *Portability* is limited by the portability of Java-enabled clients, and Java-enabled clients are ubiquitous; they exist in virtually any software environment. *Software evolution* is trivial. Once the new version of an I^2C service is available, it can be directly accessed from the network. The issue of *modularity* with respect to I^2Cnet services and their components is also vital. I^2Cnet services are comprised of small well-defined functional modules, which may be cached and shared using web caching and prefetching techniques such as proxies and browser caches. Some of these modules correspond to libraries for managing diverse forms of multimedia data. Furthermore, I^2Cnet services are autonomous and independent in the sense that they may be introduced or discontinued anytime, without disrupting the operation of the system.

I^2C services are delivered over the WWW by I^2C servers. An I^2C server consists of a Web server that distributes the Java applet itself and an I^2Cnet daemon that communicates with the service applets in order to manage the contents of the virtual workspace. Even though the objects included in the virtual workspace are, in general, distributed over a number of I^2C servers and the Web, the user is able to manipulate them as if they were physically present in the workspace. Each virtual workspace maintains an index of Uniform Resource Locators (URLs) for the objects it contains. This fact, in combination with the paperless office metaphor that is used in the user interface of the virtual workspace, gives the element of network transparency to the system and increases its usability. The communication of an I^2Cnet service with a particular I^2C server employs well-defined protocols which ensure the transfer of minimal amounts of information. The resulting bandwidth-conserving architecture allows users to interact with certain parts of the system unaware of how other services operate. Thus, the learning cycle of the potential users is short and users can easily follow the evolution of particular services.

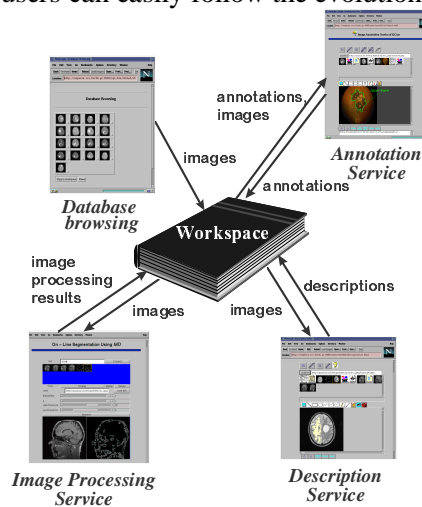


Figure 1: Virtual Workspaces provide for service interoperability, user interaction, and data persistence to geographically distributed I^2Cnet users.

3. I^2Cnet Services

3.1 Virtual Workspaces

Service integration is accomplished through virtual workspaces, which provide access to intermediate service results in addition to other forms of multimedia data. As shown in Fig. 1, the user is able to retrieve from the workspace the result of a service request and use it as input to another service. For example, the user may apply a segmentation algorithm to an image and use the segmented image as the input to an image description algorithm. The image description may then be used as input to a retrieval by content service. By logging into a workspace, the user may view the items that were inserted into the workspace during his last interaction with the workspace. Furthermore, multiple users may collaborate over a workspace, sharing material of common interest [2]. Virtual workspaces in I^2Cnet , follow the

“network computer” paradigm of information management. A network computer has minimal memory, disk storage and processor power and is designed to connect to a network, especially the Internet. The idea behind network computers is that many users who are connected to a network don't need all the computer power they get from a typical personal computer. Instead, they can rely on the power of the network servers. This is really a variation on an old idea—diskless workstations—which are computers that contain memory and a processor but no disk storage. Instead, they rely on a server to store data. Network computers take this idea one step further by also minimizing the amount of memory and processor power required by the workstation [3].

3.2 Database Browsing

I^2Cnet services allow users to browse through medical image archives guided by different notions of image similarity. These archives are maintained by medical specialists who serve as administrators/moderators of I^2Cnet sites. Thus, classes of medical images corresponding to the special needs, interests, and expertise of particular I^2Cnet sites, may be created and enriched with material which has been created locally or has been retrieved from the Web. The I^2Cnet administrator may include in the local I^2Cnet information repository image related data from the local healthcare environment, or any network-accessible place in the world. This is another example where the collection and management of content is paired with technical issues involved in system design and development, so that medical specialists share experience, information, and data regardless of their physical location.

3.3 Image Analysis and Processing Service

An extensible collection of image processing and analysis algorithms can be supported by I^2Cnet . These algorithms allow remote users to use the algorithms present at various I^2Cnet sites to analyze images from the I^2Cnet repositories or images which reside anywhere on the Web. This service is based on an *algorithm execution tool* which facilitates the dynamic introduction of new image analysis and processing algorithms. The algorithm execution tool reads an *algorithm description file* and dynamically creates the Web interface of the algorithm. This graphical user interface allows the user to select the input of the algorithm from the virtual workspace, while the output of the algorithm is automatically inserted in the same workspace. Since the user interface of the algorithm is dynamically constructed, it may take into account user-specific requirements. Hence, by augmenting the algorithm description files with user profile information, the system may be extended to provide adaptable user interfaces with respect to personal user preferences. When the algorithm execution tool receives a request to execute an algorithm, it creates an *execution agent* that takes care of all the platform specific details of the algorithm execution. This fact allows the easy introduction of new algorithms, as well as new platforms.




Figure 2: The I^2Cnet annotation format promotes the expression of different forms of annotation: comment, refutation, confirmation, correction, and illustration to name a few.

3.4 Annotation Service

The objective of the I^2Cnet image annotation service is to provide healthcare professionals with the ability not only to access medical image collections, but also to interact with imagery, creating, viewing, and communicating annotations on groups of images [4]. An image annotation may include various media types such as images, text, sound, hypertext, video as well as links to any other types of data objects (see Fig. 2). Users are able to create annotations and then use them either in a discussion forum

or keep them in a personal collection. Furthermore, a user is able to browse through annotation collections or use search techniques to retrieve appropriate material based on the annotations stored in I^2Cnet . In future versions of this service, a user will be able not only to access similar annotations to the one that has been constructed, but also to view patient record data that has been associated with the specific annotation.

3.5 Description Service

Using the I^2Cnet description service, users are able to create a description of an image and use it as a visual query in extracting similar images for the I^2Cnet database, using the retrieval by content service of I^2Cnet (section 3.6). An image description may be considered as a sentence of a visual language that describes images in terms of their visual properties. The I^2Cnet description service allows the user to import an image and specify regions of interest on it. These regions of interest may be drawn manually, or can be the output of an image segmentation algorithm which has been imported from the virtual workspace.

3.6 Retrieval by Content Service

The retrieval by content service of I^2Cnet , allows users to retrieve, from a specific image class, images with a visual information content close to that of a visual query. Various forms of visual queries are supported in I^2Cnet represented by different description types (retrieval by content algorithms). A particular description type currently available on the Web provides a query-by-example content-based query. The user interface of all the retrieval by content algorithms supported by I^2Cnet are dynamically constructed by the *algorithm execution tool* (section 3.3). The product of a retrieval by content service request is a query object which includes the query executed and its result. This query object may be used to cache query results and, thus, avoid re-execution of the query if the database has not been modified since the last execution. Furthermore, query objects may include user-feedback provided in response to query results. Hence, the system may gain information on the perspective of each user, as well as general information on how each retrieval algorithm works.

3.7 On-line Collaboration

Most I^2Cnet services facilitate two forms of user interaction: e-mail, using the mail option integrated into advanced Web browsers, and cooperation through concurrent viewing of the contents of a workspace shared by multiple users. The on-line cooperation service of I^2Cnet satisfies the need for a more direct form of communication. Based on this service, I^2Cnet users sharing a workspace may discuss using an on-line “talk” facility and comment on data objects present in the virtual workspace. In the course of an on-line session, the annotation of a group of images may be performed in collaboration. Even though, currently, the users of the service may collaborate only on images, in the future, collaboration editors will be provided for other media types as well. A new instance of the collaboration server is launched, when a user requests the creation of a collaboration session bound to specific virtual workspace. 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 exchange of information passes through the server, enabling the recording of the complete collaboration session. Furthermore, specific snapshots of the session may be taken at no additional cost and stored in the I^2Cnet annotation format. The novel aspect of this service is that, once it is completed, it will allow I^2Cnet users to discuss on-line a wide range of data types, ranging from images to patient record segments.

4. Example User Sessions in I^2Cnet

Typical I^2Cnet user sessions have several phases, i.e. the collection of data (images, annotations, etc.), their spatial arrangement in the workspace, the processing of data objects in the workspace, the communication of workspace items to other users, and the persistent storage of the virtual workspace. As a first example, consider the process of authoring an annotation in I^2Cnet . In the process of authoring an annotation, user interaction with the virtual workspace goes through *data collection*, *organization*, *authoring*, and *communication*. When an I^2Cnet user logs-into a virtual workspace, the data collection saved at the end of the previous session appears. First, the user may collect additional relevant data. After populating the virtual workspace, the user usually visits the annotation authoring page. Images

may be imported to the canvas of the annotation and overlays may be created using the available graphical tools. Several items may be selected from the virtual workspace and be linked to the current annotation. Then, the user may wish to write an accompanying report to complete the annotation. At any point in this process, the user may decide that additional data is necessary, add them to the workspace, and continue editing the annotation. In the communication phase, the user may e-mail the just-created annotation to another user, using the built-in mail functionality of the browser, or discuss it with another user using the on-line collaboration service. Alternatively, the user may save the annotation in the current workspace, or post it to *I²Cnet*. All *I²Cnet* postings are moderated by *I²Cnet* administrators. Selective postings from guest users and postings made by authorized users are linked to the *I²Cnet* database and are subsequently available to all other users.

As another example, consider a retrieval by content scenario. In this case, the user creates an image content description to use it as a visual query. In this process, the user may use several image processing services to analyze different aspects of the image. Then, using the image content description service and the data objects collected in the workspace, the user constructs an image description and provides it as a visual query to an image retrieval by content service. At any point, the query object may be saved for later reference or e-mailed to other users. For each of the retrieved images, associated annotations which present additional medical information or patient record data may be accessed. The user may use e-mail to communicate the query results to his peers, or start an on-line collaboration session to discuss aspects of the retrieved images and the associated data. The result of the discussion may be stored as an annotation that other participants may refer to. Finally, the user may use the data objects created during the on-line collaboration session to extract similar or related annotations from the *I²Cnet* database.

5. Conclusions

One of the reasons why the world-wide Web is enjoying such rapid growth, since its invention in 1993, is its loosely-coupled network-centric architecture [5]. Content information is available to all those with the right access permissions regardless of their geographic location. Access to specific information services can be accomplished via direct access to the source or via a broker or a directory server, which binds the service requester to the source(s) of information. The architecture of *I²Cnet* follows the same network-centric, service-oriented architectural paradigm. *I²Cnet* provides services that help users interact with image collections through browsing, navigation, and analysis, as well as through the addition of annotations and descriptions. Hence, not only the interaction of a large number of users with the *I²Cnet* repositories is facilitated, but also the active participation of users in the evolution of *I²Cnet* in terms of content and services is promoted.

References

- [1] Stelios Orphanoudakis, Catherine Chronaki and Despoina Vamvaka: *I²Cnet*: Content-based Similarity Search in Geographically Distributed Repositories of Medical Images, *Computerized Medical Imaging and Graphics*, Vol. 20, No. 4, pp. 193-207, 1996.
- [2] Catherine Chronaki, Manolis Tsiknakis, Stelios Orphanoudakis: Virtual Workspaces in *I²Cnet*, Presented in Mednet96, Oct 14-17, 1996, Brighton, England
- [3] William S. Rukeyser. Network Computer just for the Internet. CNN Correspondence, URL: <http://cnnfn.com/mymoney/9602/28/network.computer/>
- [4] Catherine Chronaki, Xenophon Zabulis, Stelios Orphanoudakis: Maintaining Image Annotations in *I²Cnet*. In Proceedings of EuroPACS'96, pp. 141-145, Oct. 3-5, ICS-FORTH, Crete, Greece.
- [5] Tim Berners-Lee: WWW: Past, Present, and Future, *IEEE Computer*, October 1996, pages 69-77.