S. Sengupta, P.D. Clayton

Review Paper

Columbia-Presbyterian Medical Center,
New York NY, USA

Clinical workstations: An architectural perspective

Abstract: The role of a clinical workstation is examined as an integral part of a larger, clinical information delivery and acquisition system. Different care scenarios and environmental factors influence the behavior of a workstation. The common functional components of a workstation are information resources, application logic and presentation. A workstation is successful when each of its components operates within an information architecture and contributes to meet user needs. New technologies to integrate and display information are making the workstation functions independent of the actual hardware and software platform.

Keywords: Clinical Workstation, Information Architecture, Integration

1. Introduction

A clinical workstation is usually described as a combination of hardware and software which interacts directly with a user who is interested in information related to clinical care of patients. The primary component of a workstation responsible for user interaction is the user interface, which is often graphical in nature. The information related to clinical care may reside entirely on the workstation. In today's computing environment, however, it is more likely to be distributed over a set of back-end hosts and databases, all connected to the workstation by a common network. This networked environment raises the integration aspect: the distributed nature of the information must be hidden, and a coherent, comprehensive view must be presented to the user. This integration is accomplished in stages by all participating systems on the network. A workstation is also the primary conduit of information flow between the user and the larger clinical information system represented by many systems on the network.

General-purpose workstations and personal computers have existed for over 10 years. Yet it is difficult to find implementations of integrated clinical workstations in health-care workplaces that are widely agreed to be successful. There are many reasons for this failure, the primary of which is not understanding the value of integration and its impact on tasks of the users. Experience shows that partial solutions in this area have not been shrink-wrapped and shared among institutions and users; health-care environments seem to be unique. Furthermore, a somewhat nebulous definition of the tasks of a workstation contributes to its unfinished nature: there are simply too many information sources that may be required to be integrated and then presented in too many different ways in too many different work scenarios; creation of an all-capable workstation is a never-ending quest. In the past, a discussion on clinical workstations had often resulted in enumerating relative pros and cons of individual vendor-supported hardware and (system) software platforms. Examples include Apple Macintosh,

IBM/Microsoft DOS, IBM OS/2, Microsoft Windows, etc. Today, at the front-end, newer technologies such as World-Wide Web browsers promise hardware and operating system independence and superior user choices. Similarly, at the back-end, the use of object-component technology promises complete network transparency and good interoperability. Consequently, workstation issues today are tightly related to overall software architecture and less so to issues regarding a specific hardware platform.

The rest of the paper is organized as follows: The next section explores requirements of an integrated clinical workstation, a topic well represented in the literature. First, we discuss user needs and follow-up with architectural prerequisites necessary to make a clinical workstation adequately meet the articulated user needs. Subsequently, we model a clinical workstation using distinct functional components, each of which fulfills a practical integration need. We then present examples of technologies and solutions that represent these components, and conclude with some thoughts about

possible trends in the future.

2. Requirements

The requirements of an integrated clinical workstation vary based on several criteria. One way to classify them is to consider user role, user specialty, user workplace, and the data of interest in any given scenario. A surgery (specialty) resident (role) who reviews a set of clinical results (data) about all his/her post-operative patients in a nursing station (workplace) has a different set of workstation requirements than a radiologist (specialty, role : physician) who reviews a CT-scan (data) in a clinic (workplace). A diabetic (specialty : medicine) patient (role) at home (workplace) browsing quality-of-life information (data) has a different set of requirements than a visiting nurse (role, specialty : medicine) reviewing medications (data) of a patient with tuberculosis at home (workplace). The clinical information required in each case is a different subset of the total clinical information available for patients, and the presentation content and medium are also different. These classification criteria also influence infrastructure requirements such as access control, security, and acceptable response levels.

Obviously, there is no single workstation solution that addresses all these scenarios. Solutions that address individual scenarios often look and behave differently. Partly, the underlying platform contributes to the problem: a hand-held terminal connected through a wireless communication link is incapable of providing the same workstation functionality provided by a high-powered desktop workstation connected through a high-speed local-area network. Often, workstation developments are not coordinated, or solutions are not purchased from different vendors. This situation contributes to the mismatch problem even

when the underlying platforms are homogenous in nature. When the same user is subjected to multiple presentations and behaviors, frustration is likely, and it leads to refusal to use.

A set of common user needs should be well understood in the context of information delivery to the users [1,2]. These form the basis for technical requirements for workstations. It is also extremely important to understand the prerequisites required within the information architecture that would help to meet the requirements. Even as we acknowledge the futility of a single, all-encompassing solution, we can lay down ground rules that permit integrated functional behavior in all scenarios, limited only by technological, and perhaps economic, considerations.

2.1. User Needs

1. Users expect clinical information to be available when and where they need it. Unavailability of necessary information impacts quality of care directly and negatively. Users want the information instantaneously, and it should be in the form appropriate for its need.
2. Information should be integrated regardless of how and where it is generated. Data should be logically related, accurate, and complete. All data should be available for review including data in different media (coded, narrative, image, audio, video).
3. Information access should be simple and intuitive. Information should also be made secure and be available on a "need-to-know" basis.
4. Information should link transparently with other knowledge sources. Examples include drug databases, literature on clinical trials and their results, etc.
5. Information should help a caregiver in his/her ability to initiate and follow treatment plans, and should

provide automated interpretations, suggestions, and alerts when a clinical situation warrants such attention. Harnessing of information should lead to improvement of the process of giving care.

Users treat a workstation as a tool to address their clinical information needs. Consequently, from a user's viewpoint, any unfulfilled need implies a "broken" workstation. An information systems builder, on the other hand, identifies a workstation to be a component within a larger set of computing components that work collectively to meet the user needs. Within a health-care institution, the reality is that clinical information is generated on many different technological platforms by many different departmental organizations. An integrated system is possible only when a practical information architecture is established, and an integrated system permits construction or establishment of an integrated workstation [3, 4].

2.2. Architectural Prerequisites

1. The information infrastructure must include an advanced communication network and adequate computing resources. All information producers- and consumer-systems must be networked with high bandwidth. All clinical care areas, offices, and other user workplaces must have network connections. Network connectivity should be as ubiquitous as telephone connectivity.
2. A logical set of syntactic and semantic standards must exist for information exchange. A layer of efficient and manageable middleware must exist to facilitate information exchange. The middleware must be extensible and permit incremental changes while maintaining communication with old legacy systems.
3. A collection of back-end servers must exist which contain data inde-

pendent of any application: (a) a logically centralized clinical information database server, optimized for multimedia data storage and patient-centric data retrieval; (b) a logically centralized medical vocabulary server where all medical concepts are defined and classified; (c) a decision-support engine for automated generation of clinical interpretation, suggestions, and alerts.

4. A choice of alternative front-end workstation platforms should be made based upon different environments. Multiple presentation techniques need to be simultaneously maintained. At one level, a workstation may be 100 Mbps LAN-attached and multimedia capable; at another level, it may be a personal digital assistant with a cellular phone connection. Note that we are referring to information specialist's views of the workstation, and not the user's view.

An information-system environment that supports the prerequisites mentioned above is positioned to support a workstation as conceptualized by a user. The user is likely to be unaware of and, in fact, not required to know about most of the back-end servers and services; the workstation defines the comprehensive and only view of information for the user in a given care scenario. An information architect, on the other hand, must define the functional components of a workstation and build (or assemble) the components to create workstations using the back-end services.

3. Functional Components

From a technical viewpoint, there are three overall components of a clinical workstation: *information resources and access*, *application logic*, and *information presentation*. While the clas-

sification at this level seems elementary, further refinement clearly distinguishes the functions of each hardware and software component that form the basis for the workstation.

We consider two generic functions of a workstation: information generation and information review in the following discussion.

Information Generation

Clinical information is generated in two ways. A user manually generates data by using an application (which has the three components mentioned above). A physician on an "order entry" application generates clinical orders, a technician keys in results of a microbiology test, and a radiologist dictates a mammography report. In the second way, information is generated automatically. A laboratory chemical analyzer automatically enters a blood-test result, a blood-pressure monitor generates information periodically, and a decision-support system analyzes collective clinical information of the past to generate an alert. In the former, the presentation component of a workstation permits the user to enter data, application logic determines the validity and association of the data, and resources and access find a place to store the data. In the latter, the presentation component is not necessary (other than one associated with the monitoring of an automated process).

Information Review

Review of clinical information is far more complicated. All previously generated and stored clinical information must be made available from the resources, complex application logic must relate to the scenario under which data are being reviewed, and then the presentation must match the data with the user needs in the given scenario. When review is transparently merged with generation of information (such as an order entry by a physician as a result of clinical data review, followed

by automated discovery of the order being contra-indicated, and the subsequent notification of the alert), and the entire episode occurs within seconds of a manual action, the articulated user needs are fulfilled by the workstation.

3.1. Information Resources

Ancillary systems (Laboratory, Radiology, Pathology, etc.) generate and store clinical information in hospital scenarios. It is conceivable, but highly improbable, that all clinical information generators are part of a large, monolithic application on a homogeneous hardware and software platform. A workstation must have access to integrated information; there are several ways to accomplish this. It can collect information on-the-fly from different generators when a review action is initiated; this method is time-intensive. In this case, the workstation initiates communication with ancillary systems because a review may occur randomly at any time. Or, it can collect information as and when it is generated and use it in the future; this method is space-intensive. In this case, the ancillary systems gratuitously send information to the workstation.

A practical compromise is to establish a central, patient-centric database which collects information from the ancillary systems as and when it is generated, and the workstation collects information from the database when a review occurs. The central database is the repository discussed as part of the architectural prerequisites. An additional benefit of the repository is the capability to support data-driven automated decision support [5]. As information enters the database, an event monitor checks for rules associated with the data being stored and, if found, rules are invoked.

Since rules essentially mimic human behavior, they need to have access to all clinical information about the patient in order to logically and

correctly deduce clinical interpretations, suggestions or alerts. When a user on the workstation adds, for example, a treatment plan for the patient, the automated rules are able to immediately work out the consequences based upon the latest clinical data and respond accordingly to the user. Note that the centralized nature of the database being discussed is logical [6]. In reality, the database, under the covers, may be physically distributed. There are several possible strategies. When a specific patient is being reviewed at a workstation, a replica of all of his/her full clinical information may be made available on the workstation. Or, several databases may be established, each with a non-overlapping set of patients. Each architectural strategy affects the way the workstation functions are to be designed and implemented.

Secondary to the clinical information, a workstation needs to access reference material. For scholarly material, Medline and other clinical advisories are popular. For a nurse or an administrator, there are administrative tasks related to clinical information which, when automated, provide added incentive to use workstations. Clearly, these resources and applications must also be available on the underlying information infrastructure. These are also examples of systems that are unlikely to be part of a monolithic and homogeneous clinical information system but are important to the workstation from the user's perspective.

3.2. Application Logic

The list of applications on a clinical workstation is too large to be enumerated. Clearly, some applications are common across user scenarios, and some are specific to the task at hand or the environment. We broadly identify the set of applications into two classes: *technical* and *clinical*. Technical applications or modules support the primary, clinical applications. Clinical

applications accomplish a necessary health-care task based upon the environment.

Technical applications can be further classified as system-management functions and utilities. A critical system-management subsystem is the security subsystem, which identifies and authenticates users, maintains an audit log, performs authorization checks, determines if data need to be encrypted, supports automatic sign-off depending upon an environment or scenario and, in general, watches for attacks and violations. Security functions, associated with all interactions between the system and a user, must be integrated with clinical applications, but it is a separate collection of applications, perhaps acquired from a vendor not specialized in clinical applications. Other examples of system-management application include interface engines, performance-measurement tools, application installers, enterprise-wide electronic mail systems and hierarchical storage and backup systems. These systems are employed institution-wide, and are typically distributed over several computers, helping to provide faultless and efficient operation of information exchange between information resources and users.

Utilities are personal applications helping a user in smaller tasks, and are customizable to satisfy a user's personal preferences. Examples include a word-processing application that a radiologist may invoke to update a radiology report, a multi-media viewer application that shows audio and video segments, a personal scheduler, an electronic mail viewer, etc. In practice, these utilities are mostly applications that reside and run on the workstation itself (although they may be stored on servers on the network for effective control, as determined by the application installer, a systems management application).

The users interact primarily with

clinical applications. Here, clinical data are read and written from the information resources (such as the patient repository), and analyzed and modified to fit the desired care environment. These applications model the different requirements based upon the care scenarios. For example, a clinical data review application for a surgery resident extracts clinical summaries for only the set of patients in that resident's service of interest from the very recent set of results (past 2 days). On the other hand, an application that creates a review of a single patient for a primary-care physician in the outpatient scenario collects detailed, long-term information about the patient. In contrast, the focus of review of an ICU patient is short-term data. Clinical applications encapsulate clinical knowledge required to address a given care scenario. They may use clinical systems tools, such as a controlled medical vocabulary [7], to consolidate information from heterogeneous sources, and present a coherent view to the users.

The important consideration, however, is how technical and clinical applications are integrated transparently. Indeed, integration of these functional components, perhaps more than the presentation-level components, determines the success of the workstation in meeting user needs. The applications must make themselves flexible and open so that various components may be defined and customized to work together as a single application. This is not the case today, but new directions in component technology and new standards for defining interactions between applications are promising to achieve this goal.

3.3. Presentation

The presentation function is very important to a workstation because it determines the ease of user interaction with the overall system. A very ca-

pable set of back-end resources and very sophisticated application logic will fail if the presentation is not intuitive, simple, and appropriate for the task at hand. In the past, the presentation component was often considered the most important part of the workstation with somewhat naive claims such as: "by definition, a workstation with graphical user interface is better than a workstation with simple, character-based displays." Today, the capabilities of a workstation are measured by overall functionality of the system, including the appropriateness of a user interface that matches its intended functions and scenarios. It is also clear that visual dexterity alone cannot substitute for lack of good contents and intelligence that one finds with robust information resources and mature applications. Alternate interaction media, other than character-based and graphical, such as voice-based and pen-based, demonstrate that different interfaces suit different scenarios. During surgery, and not during communication with a patient, a voice-based interface may be appropriate. During rounds, a pen-based interface may be preferable, and in the office, a multi-media, graphical and voice-based interface may be preferable. In communication over very slow-speed networks, such as the wireless networks today, a character-based interface works the best. More importantly, the design of the presentation should focus on how quickly and intuitively a task may be completed, instead of visually stimulating and with seldom necessary functions or displays.

Presentation function must use prevalent health-care metaphors and emulate existing health-care processes and practices. Users, who are less intimidated by technology today, need to find familiar emblems in the user interface, which makes the interface more intuitive. Furthermore, in order to personalize the interface, a user's preference must be taken into account

and used if the environment accepts and permits such customization. For example, all other things being equal, one user may choose to use Microsoft Word, and the other may choose to use WordPerfect for word processing functions, and the user interface must support simple ways to specify and fulfill these preferences. For seamless integration among applications, it is desirable that the presentation associated with all these applications be consistent. Today, this level of consistency is hard to achieve where applications come from different vendors, and on different computing and presentation platforms. In the past, it was hard to even interact with an application due to the heterogeneity problem. For example, if the workstation platform was Unix with X-Windows, it was impossible to interact with a Microsoft Windows-based program needing data from Novell Netware servers. Today, these problems are being solved by the vendor community: A Microsoft Windows program is made available through Windows emulation, or by using techniques that distribute the Microsoft Windows' graphical user interface over X-Windows from an application server.

4. Discussion

Clinical information systems initially started as monolithic systems: a single computer with a special-purpose database, and a large amount of application code analyzing and presenting data through directly connected terminals. These terminals also were specific to the vendor who built the computers: IBM 3270 for IBM System 370 and 390 mainframes, DEC VT100, VT220, and VT320 for DEC VAX 750 and 780, and DG Dasher for DG MV 15000, etc. Heterogeneity implied total isolation, including separate networking strategies. Information resource, application logic, and

presentation, all were completely and inseparably part of a single application, leading to duplication of every resource right down to two different terminals if one needed to access two applications on different platforms. The advent of general-purpose Local-Area Networks based on Ethernet, and generic computers as servers and clients, and the eventual acceptance of TCP/IP as the communication protocol of choice has made the environment far more integrated. Solutions based on distributed systems, including powerful system-management applications and utilities, have become increasingly available. The presentation-hardware technology has blossomed into faster, lighter, and more colorful displays, integrated audio capabilities, and video input devices (camera) as well as real-time digital video players. The networking technology routinely delivers 100 Mbps bandwidth and promises communications in excess of 1 Gb/s speed in the near future. Information-resource storage and access have become more efficient due to practical relational databases. Far more sophisticated analysis and visual displays are possible today because of the tremendous processing power available on computers. Enterprise-wide systems have evolved to be a collection of back-end servers and a large number of front-end user workstations, all communicating over high-speed networks and collaborating to present a single view of the collected information. With this technology, it is possible to build substantial functionality in clinical workstations that meet user needs to a greater degree.

Several new enabling technologies will bring the next level of integration. Two examples of such technology are: World-Wide Web (Web) and Distributed objects. The Web brings the first platform-independent graphical display capability. Also, through helper application technology, it brings a

model of seamless invocation of application logic. Finally, it brings techniques to collect information from resources located throughout the world, thus improving the quality and quantity of information to be presented to the user. Java technology offers solutions for effective and controlled distribution of application logic. Web-based presentation systems, in conjunction with Java technology, promise to be the next generation of user-interface development for the health-care industry [8,9].

Distributed object technology, represented by Microsoft Distributed Object Linking and Embedding, and Object Management Group's Common Object Request Broker Architecture [10] and OpenDoc technologies, promise to address seamless integration at the application-logic level. They offer a flexible and extensible architecture whereby new functionality can be added or inherited to change behavior of existing systems. Furthermore, when integrated with a Distributed Computing Environment [11] with its security and directory services, and with the Web, object technology creates a very powerful infrastructure to address all articulated user needs and architectural prerequisites for clinical workstations.

5. Conclusions

The experience of the past decade has shown that a clinical workstation is an integral part of a larger, clinical information-delivery system. Workstations are expected to perform different tasks based upon clinical scenarios and other environmental factors. Thus, there are several instances of a workstation, all of which fulfill a set of common user needs and can be viewed as having a common set of functional components. Each component behaves differently based upon the scenario, hence creating a different instance. The progress of computing technology, along with several positive trends, point towards achieving better integration and improved capabilities at the workstation.

References

1. The Health Care Professional Workstation, Eds: Llauro JG, Hasman A. *Int J Biomed Comput (Special Issue)* 1994;34.
2. Safran C. Defining clinical workstation [editorial]. *MD Comput* 1993;10:145-6, 192.
3. Van Mulligen EM, Timmers T, van Bommel JH. A new architecture for integration of heterogeneous software components. *Meth Inform Med* 1993;32:292-301.
4. Clayton PD, Sideli RV, Sengupta S. Open architecture and integrated information at

Columbia-Presbyterian Medical Center. *MD Comput* 1992; 9:297-303.

5. Clayton PD, Hripcsak G. Decision support in healthcare. *Int J Biomed Comput* 1995;39:59-66.
6. Johnson SB, Friedman C, Cimino JJ, Hripcsak G, Clayton PD. Conceptual Data Model for a Central Patient Database. In: Clayton PD, ed. *Proceedings of the Fifteenth Symposium on Computer Applications in Medical Care*. New York: McGraw-Hill Inc, 1991: 381-5.
7. Cimino JJ, Clayton PD, Hripcsak G, Johnson SB: Knowledge-based Approaches to the Maintenance of a Large Controlled Medical Terminology. *J Am Med Informatics Assoc* 1994;1:35-50.
8. Cimino JJ, Socratous SA, Clayton PD. Internet as clinical information system: application development using the World Wide Web. *J Am Med Informatics Assoc* 1995;2:273-284.
9. Kohane IS, Greenspun P, Fackler J, Szolovits P. Accessing Pediatric Electronic Medical Record Systems via the World Wide Web. *Pediatr Res* 1995;37:139A.
10. CORBAMED. <http://www.omg.org/corbamed/home.htm>.
11. DCE and Objects, by Dave Chappell, Chappell and Associates. <http://www.osf.org/dce/3rd-party/ChapRpt1.html>.

Address of the authors:

Soumitra Sengupta, Paul D Clayton, Columbia-Presbyterian Medical Center, Atchley Pavilion Rm 1301, 161 Fort Washington Avenue, New York, NY 10027, USA.