Architecture of an Image Capable, Web-Based, Electronic Medical Record

Steve G. Langer

With each medical center department creating and maintaining its own patient care–related data, nursing and house staff may find it confusing to log into all the information systems necessary to achieve a global perspective of the patient's state. The Medical Information Network Database application provides a logically centralized Worldwide Web viewing application for the physically distributed data. In addition to coordinating data displays for histories, laboratories, pathology, radiology, and discharge summaries, the application can be configured to apply rule sets to the data and remind caregivers of follow-up tests or of possible reactions to treatment protocols. The viewing client runs on any HTML 2.0-compliant browser, although certain applet enhancements (notably for viewing radiological images) require a browser with Java abilities. With this "thin client" approach, the application can be configured to coexist with other applications (such as a PACS viewer), thus centralizing information and reducing the overall number of computers in the medical center.

KEY WORDS: Worldwide Web, electronic medical record, DICOM, clinical information system.

THE DATA associated with a patient's hospital stay is voluminous. At the point of admission, the hospital information system (HIS), scheduling systems (SS), and the associated admission/discharge/transfer (ADT) systems must record the relevant patient demographics, the bed the patient will occupy, and all subsequent events of care. In addition, the HIS (or the responsible departmental information systems) must record particular examination times, results, and billable information. From the care provider's viewpoint, this information avalanche can be overwhelming. Furthermore, much of it is unrelated to the immediate patient care. A successful electronic medical record (EMR) should benefit providers by filtering the information, providing only what is necessary, and also performing elementary analysis if desired.

Typically, providers are more concerned with the patient's immediate medical information. Even given this subset, with each department creating and maintaining its own patient care–related data, nursing and house staff may find it confusing to log into all the information systems necessary to achieve a global perspective of the patient's state. Figure 1 shows the user's viewpoint of a distributed data, distributed interface, clinical information systems (CIS). Note that for each department, the user may need a different computer, using a different operating system, with a login and password that differ from those of any other system. Combined with the likelihood that not all workstations are conveniently located, such systems at best force caregivers to search for all relevant workstations, and worse, may preclude the caregiver from finding all available information.

There are commercial and home grown systems that provide a centralized view of the distributed patient data, but such systems often are based on expensive platforms and require specialized (UNIX or X-terminal) workstations for the end users (Hewlett-Packard CareVue, Palo Alto, CA; Eclipsys Sunrise Clinical Management, Atlanta, GA). This is especially true in large institutions, where the total cost of ownership (TCO) favors centrally administered terminals over an EMR based on hundreds of Microsoft Windows client workstations (Microsoft, Redmond, WA). Such systems typically use a large central computer to receive data from the departmental servers, and then distribute the results to the workstations as shown in Fig 2. Although a considerable improvement over Fig 1 (the user needs to learn only one interface and login/password), costs usually constrain the numbers of workstations so providers may not find one nearby, or it may be in use by someone else. Also, reliance on a large central database computer often does not scale well when the number of supported workstations climbs beyond a few hundred, and the medical center is tied to a single vendor. Furthermore, the point-of-care workstations often are dedicated to the CIS application only. This leads to

From the Department of Radiology, University of Washington, Seattle, WA.

Supported in part by a National Library of Medicine grant (N01-LM-6-3545) and a University of Washington Integrated Advanced Information Management System grant (G081M05620-01).

Address reprint requests to Steve G. Langer, PhD, Department of Radiology, University of Washington, Box 357115, Seattle, WA 98195.

0897-1889/00/1302-0006$10.00/0
doi:10.1053/di.2000.5424

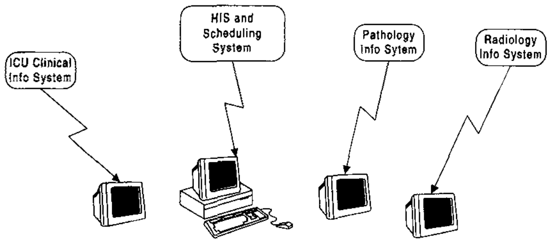

Fig 1. A distributed data, distributed interface (DDDI) approach to clinical information systems.

the profusion of computers and monitors often seen in medical centers, and increases networking, power, cooling, and maintenance costs.

The technologies of the Worldwide Web (WWW) offer a more portable and cost-effective approach, and would potentially allow physicians to investigate patient status from any networked computer as long as it was equipped with a Web browser. This could include office and home systems, and thus reverse the growing trend of overspecialized computer systems. Several investigators have taken steps to move toward an integrated WWW-based EMR. However, to date these efforts have largely been applied to a few specific departments.[1-4] Our institution's Medical Information Network Database (MIND) provides a logically centralized WWW interface for all available physically distributed data. In addition to coordinating data displays for histories, laboratory, pathology, radiology, and discharge summaries, the MIND database can be configured to apply rule sets to the data and remind providers of follow-up tests or of possible reactions to treatment protocols. The MIND viewing client runs on any HTML 2.0-compliant browser, although certain "applet" enhancements (notably for viewing radiological images) require Java-enabled browsers. With this "thin client" approach, MIND can be configured to coexist with other applications, thus reducing the overall number of computers in the medical center.

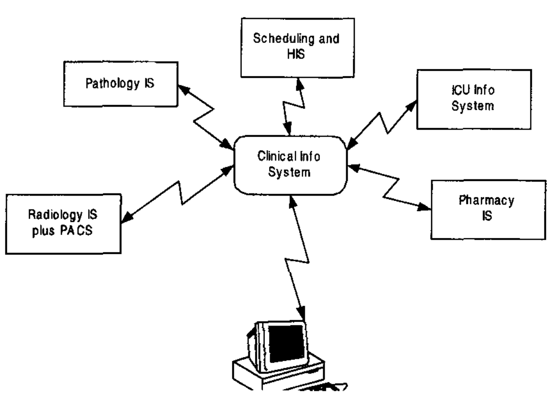

Fig 2. A distributed data, centralized interface, centralized distribution (DDCICD) approach to clinical information systems.

There are a variety of data flow scenarios for implementing a hospital wide, Web-based EMR. The simplest (shown in Fig 3) assumes that there exists a single HIS/SS that is used to schedule all patient services.[5] For each new patient event, the SS component queries the departmental information system (IS) for available examination times. Furthermore, the HIS maintains only the status of patient examinations, not the actual examination results. Thus, when a user queries the HIS for data on a particular patient, the system interrogates and returns pointers to the computers that contain completed examination data. The user may click on these pointers to call up the information directly from the departmental IS.

There are many advantages to the distributed data, central interface, distributed distribution (DDCIDD) method. Every department can manage its data as it sees fit, duplication and synchronization of parallel databases is avoided, and network usage is distributed over the entire medical center's network rather than concentrated on the subnet containing the HIS. There also are some difficulties. For instance, each IS must export HTML (Hypertext Markup Language) directly to the viewing browser, or the browser and IS must share some other common communication protocol. Furthermore, this approach is not possible for medical

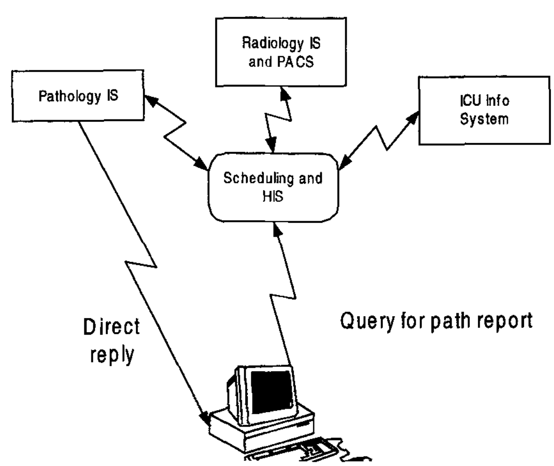

Fig 3. The DDCIDD model. This model assumes that the HIS has total knowledge of all examinations scheduled for every patient, or alternately can interrogate all the departmental ISs in real time.

centers that have a nonexistent or incomplete HIS infrastructure, that have saturated network service precluding real-time examination queries, or that deal with a large trauma population for which preexamination scheduling is unlikely. To some degree, the latter two elements are the case at our center, where we have adopted a DDCICD EMR. With this system, every departmental IS assumes responsibility for uploading patient examination data to the central archive via a Health Level 7 (HL7) network link. At times such uploads are nearly instantaneous; at others they may be delayed by 10 minutes or more. Figure 4 shows this model.

Figure 4 is topologically similar to Fig 2, but rather than a platform-restrictive central server, the distribution unit is a Web server. While the DDCICD model suffers from some of the issues mentioned previously (in particular it concentrates network traffic on the EMR server subnet), it offers potential advantages in the application of centralized rules-based patient monitoring. In addition, because all data are processed by the central system, the user interface can be made to appear identical regardless of which IS the data are coming from. This would likely not be the case if each department's IS exported data directly to the client.

MATERIALS AND METHODS

Up to this point, the discussion has focused primarily on the overall architecture of MIND. As such, Fig 5 is a considerable simplification of the overall data flow within the institution. Before the multiple HL7 inputs to the MIND cluster, an HL7 interface engine (Cloverleaf, PCS/Healthcare, Cincinnati, OH) is used to separate billing, inventory, ADT (admission/discharge/transfer) and other data to their respective destinations.

As currently designed, the MIND application consists of a multivendor UNIX cluster. IBM RS6000s act as HL7 to structured query language (SQL) gateways (IBM Corp, Westchester, NY). The gateways map incoming data from the departmental IS system's HL7 segments to tables in the SQL database. The actual database is maintained by a cluster of Suns running Ingres SQL (SUN, Palo Alto, CA; Computer Associate's Ingres, Islandia, NY). Web functions are performed by Sequent servers (Sequent Corp, Beaverton, OR). MIND queries from client browsers to the database are conducted via client-side Java applets that send their queries to a server-side common gateway interface (CGI) code, which makes the actual SQL request. The database returns data to a formatting process that transmits the representation to the browser. Figure 5 demonstrates the data flow.

In addition to textual information, we have integrated radiological image display capability with the MIND client.[6,7] This is accomplished by passing the patient's examination number (the RIS accession number) from within the MIND view as arguments on the uniform resource locator (URL) call to the client-side viewing applet. The viewing applet is then started; it automatically queries the PACS web server for the unique study by forwarding the accession number, and loads the images (Medweb, San Francisco, CA). This client is fully DICOM capable, and supports decoding of wavelet-compressed images with up to a 16-bit window and level capability. The full range of image objects is supported (digital radiography, digital fluoroscopy, MRI, CT, ultrasound, and nuclear medicine). The compres-

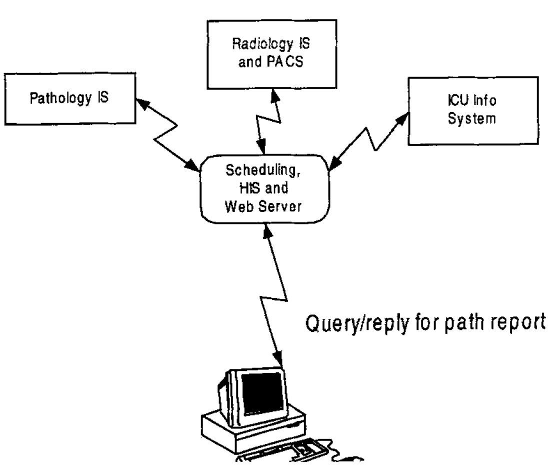

Fig 4. The DDCICD model. This is the model actually used by MIND.

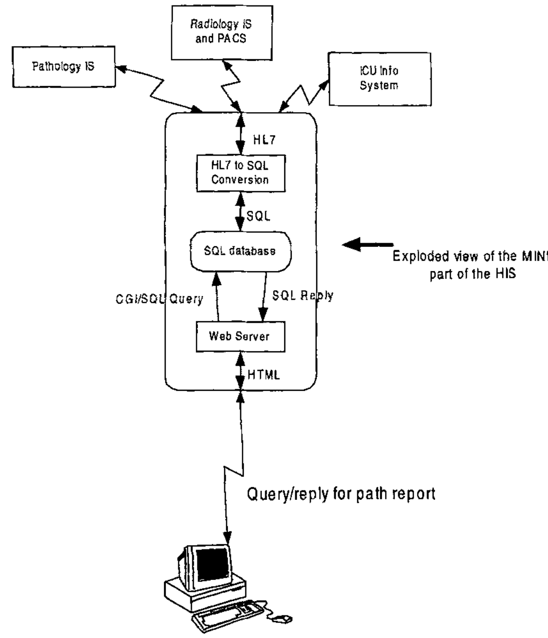

Fig 5. The steps involved in a MIND query. This is essentially an exploded view of the SS/HIS and Web-server cluster shown in the Fig 4. Note that bidirectional communication between the client and MIND database is possible, permitting data retrieval and entry at the point of care. However, for point-of-care entered information to be back-propagated to the department IS, there must be a bidirectional HL7 link between the MIND database and the departmental ISs, as well as an interface broker component to the gateways to permit SQL to HL7 writes back to the ISs.

sion level is user-selected before study download to any of the following (subjective) quality levels: lowest, low, high, and highest. The viewing client performs window/level, cine', flip, rotate, invert gray scale, zoom, annotate, pan, save image as JPEG, and other operations. Invocation of the image viewing applet is transparent to the MIND user, unless the browser does not have the Medweb applet, in which case the user must retrieve the applet from a secure server. Viewing clients are available for Microsoft Windows 95 (and newer) operating systems as well as for the Macintosh and Linux computers.

Security issues are dealt with via a trusted validation process. For example, upon initial login, the user provides the username and password to the Web page and the password is hashed using embedded algorithms (RSA, San Mateo, CA) commonly found on most Web browsers. Assuming that the hashed password (clear-text passwords are not transmitted across the network) and username match a valid account, the validation process supplies a "cookie" (containing a private key) to both the MIND server and the workstation. At this point a secure sockets layer (SSL) based encrypted session begins, and patient accesses are logged via the session cookie, thereby providing an audit trail of the patients accessed by the provider.

The graphical user interface for the patient data view was prototyped after a review of existing Web interfaces to medical records applications.[8,9] Afterward, a panel of representative physicians assisted in an iterative development process.[10] A tabbed folder metaphor was used to organize patient data. The user need only click on the tabs to view the following data: demographics, problems, laboratory results, radiology procedures, medications, allergies, surgical procedures, etc. Online references also were made available such as MEDLINE, drug reference information, laboratory references, and clinical guidelines. From initial user responses, it became clear that response time, presentation of procedure reports, and available reference data required improvement. These issues were addressed by deploying the Web service on a parallel cluster of computers, standardizing the appearance and format of transcribed reports, and including more reference links, respectively.

RESULTS

In its current configuration, the MIND database permits the following views of patient-related data: demographic information, ICD-9–based problem lists, medication lists, immunization information, allergies, clinic visit dates, hospitalization dates, procedure reports, laboratory results, radiology transcription, and images. It also supports a "clinical reminders outcome system," which identifies patients with specific clinical or demographic conditions that suggest preventative maneuvers. For example, reminders may include disease-specific processes, a physiological outcome measure (glycosolated hemoglobin), or a functional outcome measure (eg, SF36).

There are many possible data viewing formats within MIND. For instance, Figs 6-9 show some of the various "patient-centric" views available for a given patient within MIND. A patient's record may be retrieved via name, medical record number, or partial name match. For research, records may be searched by specific findings or diagnoses. Also, a provider-centric view is possible that lists all patients under the care of a given physician. For this study (and a demonstration site), a pseudo-patient was constructed.[11] In the live application, the tabs across the top of the page and links within the documents would be used for navigation. Selecting these invokes processes on the MIND server that request and format the desired data.

In addition to patient related information, MIND provides centralized access to many kinds of medical information. For instance, a provider may note that a given patient's glucose level is abnormally low. They may ask "Is this caused by one of the medications?" With MIND, the answer to the question can be found by simply clicking on the drugs displayed under the patient's medical record. Relevant information on the given pharmaceutical is pulled from online resources such as the medical center drug formulary. Access also is provided to MEDLINE, laboratory reference information, and

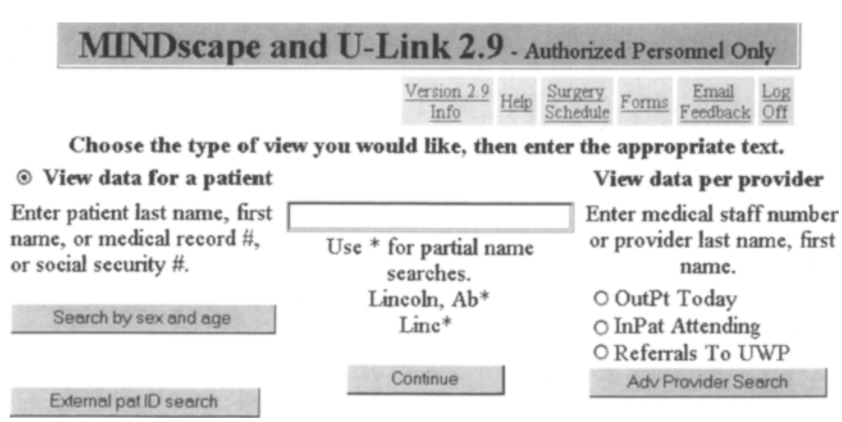

Fig 6. The MIND secure introduction and login page. The provider is only able to reach this point after supplying a valid account name and password (both are sent in cyphertext). From this point forward, activities are logged and data are sent over an encrypted channel.

Patient: PT ZZTEST U6999999 Male Age: 19	Select New Patient	Help	Email Feedback	Log Off

Demographics Problems Meds Allergy Providers Visits Transcripts
Lab Radiology Pathology Reminders Immuniz Procedures Findings

PATIENT HAS REMINDERS DUE. see them

i Click to access University Healthlinks Care Provider Toolkit.
i Click to access University Healthlinks Resources by Clinical Specialty.
i Click to search Medline.

Problem List

Problems are recorded manually by clinic staff. To view any additional potential problems see symptoms.

ref. lib.	Problem Description click for detail	Date of Onset	Date of Episode	Nbr of Episodes	Status
i	ESSENTIAL HYPERTENSION	unreported	unreported	unreported	Active
i	FOREIGN BODY ESOPHAGUS	unreported	unreported	4	Active
i	HEADACHE THIS IS 3/23/98	unreported	unreported	unreported	Active
i	INTRACRANIAL HEM NEC/NO	unreported	unreported	unreported	Active

Fig 7. The problem page includes a rapid summary of the patient's clinical evaluation. This is the typical jumping-off point for a provider seeing a patient for the first time. Selecting either the "i" links or a given problem (eg, "headache") produces a grid of all the laboratory and other findings relevant to that problem.

Patient: PT ZZTEST U6999999 Male Age: 19	Select New Patient	Help	Email Feedback	Log Off

Demographics Problems Meds Allergy Providers Visits Transcripts
Lab Radiology Pathology Reminders Immuniz Procedures Findings

Clinical Reminders

Status	Condition	Service	Eligibility	Min Freq	Last Done	Due On
DUE*	Diabetes	Cholesterol Test	Diabetes dx: >=18	q3yr	Unknown	Unspecified
not due	Diabetes	Foot Exam	Diabetes dx: >=18	q1yr	15-dec-1997	15-dec-1998
DUE*	Diabetes	Glycated Hemoglobin	Diabetes dx: >=18	q6mo	Unknown	Unspecified
DUE*	Diabetes	Obesity Screen	Diabetes dx: >=18	q3yr	Unknown	Unspecified
DUE*	Diabetes	Retinal Exam	Diabetes dx: >=18	q1yr	Unknown	Unspecified
DUE*	Diabetes	Urine Protein	Diabetes dx: >=18	q1yr	Unknown	Unspecified

Fig 8. The reminder page. This prompts providers to perform follow-up examinations or laboratory tests.

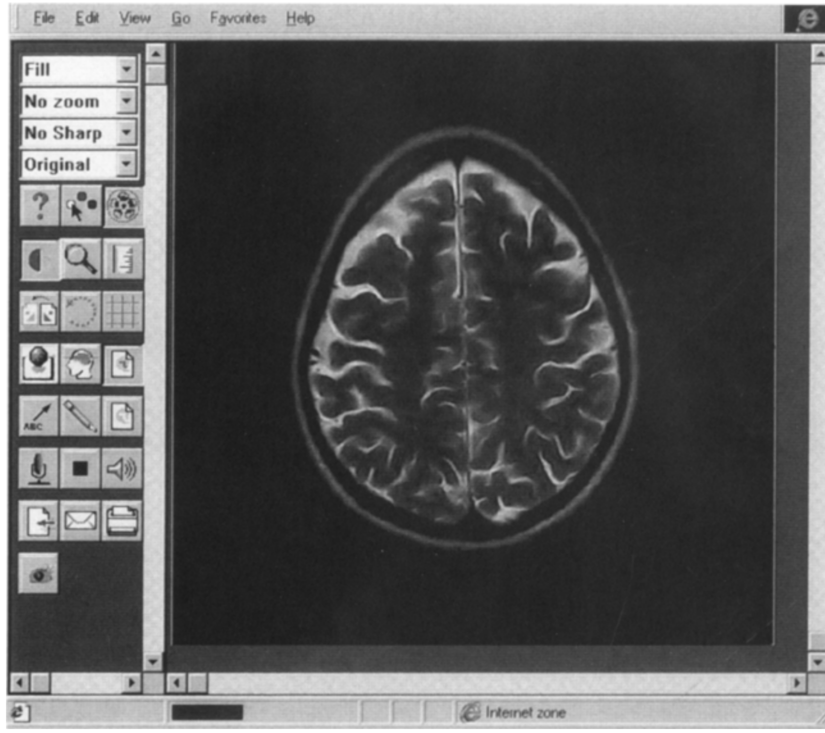

Fig 9. The DICOM viewing applet in use.

a context sensitive, critical care pathway, decision support system—an often requested feature.[12] MIND also provides tools to help the physician select medications that are covered by the patient's insurance provider. This is done by cross-checking the selected treatment drugs in the medical center formulary against the formulary of the patient's insurance company. Any prescribed drugs that are not covered by the patient's insurance are flagged, allowing the physician to make a substitution if appropriate. Because our center serves an ethnically diverse population, MIND offers (via the demographic information associated with a patient) links to cultural reference materials based on ethnicity and religion.[13]

Attempts at a comparative analysis between MIND and the pre-existing paper record system are somewhat difficult because few in-depth studies were performed on the paper-based system. Therefore, the performance metrics listed in Table 1 for the paper-based system represent the consensus obtained from several users who were experienced with it and other sources.[10,14] In contrast, the values in the MIND column represent worst-case estimates of measured performance, which varies as a function of network and server loading. (An exception is the estimate for "lost" charts. Although a MIND chart is never lost, initial typing errors in a patient's name or medical record number can make it difficult to find the record later. Typically, the problem can be circumvented by doing a partial match search within MIND on a portion of the patient name or medical record number, thus producing a list of all likely candidates. When the discrepancy is ultimately found, the record is corrected. The "lost" figure below is based on the number of times this has been necessary, but this estimate may be low if users simply have not reported the error.)

Clearly the most dramatic points in Table 1 are the accelerated turnaround of information, and in some cases (eg, insurance checks for drug reimbursements) the availability at all. Furthermore, when used in conjunction with other computer-assisted methods (eg, voice recognition), the acceleration of diagnostic reports to the primary physician may greatly increase the referrer's confidence that diagnostic examinations will provide timely input in treatment decisions without requiring "out-of-band" phone consultations with the diagnostic physician.

DISCUSSION

MIND has been in use for over 2 years. In that time it has gone from alpha testing to having more

Table 1. Comparison of Several Basic Performance Metrics: MIND Versus Paper Records

	MIND	Paper Records
Lost paper chart occurrence	<2% (typo in record)	~20%
Chart call to delivery	<10 s	>30 min
Time from radiology dictation to reach EMR (human transcription)	4 to 6 hr	Overnight
Time from radiology dictation to reach EMR (with voice recognition)	<60 s	Not tested
Time to cross-check drug payment on insurance plan	<60 s (if insurer's formulary is online)	>1 hr (if attempted)
Time to research drug side effects	<10 s	>10 min

NOTE. Refer to the text for details.

than 1,000 enrolled users. Where it has been installed, it has proven enormously popular, not only because it supplants several client computers but also because the user interface is self-explanatory to anyone with experience using Web browsers.[11] Because of its simplicity, MIND's training costs are marginal, as is the need for specialized support.

MIND represents the culmination of 3 years' work from numerous programmers, with input from dozens of providers throughout our medical center. Even so, it continues to evolve with additional inputs from users and as technology changes occur among the various departments. For instance, the radiological image client is evolving from the current Medweb platform to an integrated Web service provided by our PACS vendor. In addition, there is interest in providing a broader availability of MIND to physicians in our referring clinics. Such an implementation, however, requires a method to restrict patient access for providers to only those patients for which they are responsible. In addition, maintenance issues alone consume the time of two full-time software engineers in the medical center's information systems department.

In its current incarnation, MIND is a read-only, data display tool. Thus, there is not yet a provision for point-of-care entry of new information. Obviously, this would be a great boon for updating bedside charts, physician's notes, etc. This capability is now under development. Unfortunately, the problem is considerably more complex than initially perceived. Changes entered at the point of care must back-propagate not only to the MIND database but also to the proper individual department IS as well (if one exists). To accomplish this, the existing SQL-to-HL7 gateways will need to have a bidirectional broker component. In addition, not every departmental IS permits or has the capability to accept inbound HL7 messages. Hence, many of our legacy IS systems will need interface upgrades as well. Finally, our network will require the installation of institution-wide upgrades to maintain performance as use increases.

In addition to improving the geographical availability of clinical information, we believe MIND will increase the timeliness of it. This will come about from several synergistic effects (which are ongoing projects):

- Real-time, point-of-care data entry will reduce delayed entry of physician's notes.
- Implementation of medical center–wide voice recognition capability or structured reporting templates will shrink transcription turnaround from hours to minutes.
- Expert systems use will increase, suggesting relevant tests and reminders.

MIND is part of an evolving strategy to link all resources of the medical center into one unified medical informatics paradigm. As such, its portability, speed, ubiquity, and expert systems amplify the abilities of nursing and medical staff alike. Furthermore, the integration of medical information by patient, provider, or diagnosis/findings simplifies both research and clinical decision support tools. The future of MIND includes more decision support and even tighter integration with visual data such as cardiac wave forms. In conclusion, MIND will become increasingly central to our practice, and we believe other practitioners will benefit from pursuing similar systems.

ACKNOWLEDGMENTS

The author thanks Peter Killcommons of Medweb for PACS web-server support, and Michael Richardson, MD, for hosting the demonstration page. Additionally, the author is grateful for the ongoing support of the medical center's information systems division and the vision of the faculty in the Division of Biomedical Informatics.

REFERENCES

1. Barbaras L, Parker JA, Donohoe KJ, et al: The all-digital department moves to the Web. RSNA-EJ 1997, vol 1

2. Holman BL, Mammone GL, Greenes RA, et al: Brigham-RAD: An approach to education, decision support, information dissemination, and collaboration via the World Wide Web. RSNA-EJ, 1997, vol 1

3. Appleby C: Web-o-matic isn't automatic–yet. Internet technology hasn't broken the barrier between doctors & computers. Hosp Health Network 1:30-34, 1997

4. O'Kane KC, McColligan EE: A web access script language to support clinical application development. Comput Methods Programs, Programs Biomedicine 55:85-97, 1998

5. Langer SG, Wang J: A goal-based cost-benefit analysis for film vs. filmless radiology departments. J Digit Imaging 9:104-112, 1996

6. Stewart BK, Langer SG, Taira RK: DICOM image integration into an electronic medical record using thin viewing clients. Proc SPIE 3339:322-328, 1998

7. Stewart BK, Langer SG: Integration of DICOM images into an electronic medical record using thin viewing clients. AMIA Symposium 1998, pp 902-906

8. Kohane IS, Greenspun P, Fackler J, et al: Building national electronic medical records via the World Wide Web. JAMIA 3:191-207, 1996

9. Cimino JJ, Socratous S, Clayton PD: Internet as clinical information sy stem: Application development using the World Wide Web. JAMIA 2:273-284, 1995

10. Tarczy-Hornoch P, Kwan-Gett TS, Fouche L, et al: Meeting clinician information needs by integrating access to the medical record and knowledge resources via the Web. Proc AMIA Ann Fall Symposium 1997, pp 809-813

11. Richardson ML: MINDscape and PubMed: Web sites that can change the way we work. Acad Radiol 5:519-520, 1998; The MIND demonstration site is at *http://www.scar.rad.washington.edu/msdemo.html*

12. Smith R: What clinical information to doctors need? Br Med J 313:1062-1068, 1996

13. Ethnomed at *http://healthlinks.washington.edu/clinical/ethnomed/*

14. Chimosky SJ, Schmuland D, Boyle EM, et al: Second Opinion: Do you think electronic medical records will effect your practice? University of Washington MedAlum 22:11-13, 1999