SynEx
SYNERGY ON THE EXTRANET

providing intuitive access to federated healthcare records - securely

London Demonstrator Site

Final Report

Authors:

D. Kalra, A. Austin, A. O'Connor, D. Patterson, D. Ingram, D. Lloyd

CHIME

UCL

Contents

Deliverable Reference Information

Introduction

The key ingredients of the SynEx-UCL software components are:

1. <u>A comprehensive and federated electronic healthcare record</u> that can be used to reference or to store all of the necessary healthcare information acquired from a diverse range of clinical databases and patient-held devices.

2. <u>A directory service component</u> to provide a core persons demographic database to search for and authenticate staff users of the system and to anchor patient identification and connection to their federated healthcare record.

3. <u>A clinical record schema management tool</u> (Object Dictionary Client) that enables clinicians or engineers to define and export the data sets mapping to individual feeder systems.

4. <u>An expansible set of clinical management algorithms</u> that provide prompts to the patient or clinician to assist in the management of patient care.

CHIME has built up over a decade of experience within Europe on the requirements and information models that are needed to underpin comprehensive multi-professional electronic healthcare records. The resulting architecture models have influenced new European standards in this area, and CHIME has designed and built prototype EHCR components based on these models. The demonstrator systems described here utilise a directory service and object-oriented engineering approach, and support the secure, mobile and distributed access to federated healthcare records via web-based services.

The design and implementation of these software components has been founded on a thorough analysis of the clinical, technical and ethico-legal requirements for comprehensive EHCR systems, published through previous project deliverables and in future planned papers.

The clinical demonstrator site described in this report has provided the solid basis from which to establish "proof of concept" verification of the design approach, and a valuable opportunity to install, test and evaluate the results of the component engineering undertaken during the EC funded project. Inevitably, a number of practical implementation and deployment obstacles have been overcome through this journey, each of those having contributed to the time taken to deliver the components but also to the richness of the end products.

UCL is fortunate that the Whittington Hospital, and the department of cardiovascular medicine in particular, is committed to a long-term vision built around this work. That vision, outlined within this report, is shared by the Camden and Islington Health Authority and by many other purchaser and provider organisations in the area, and by a number of industrial parties. They are collectively determined to support the Demonstrator Site as an ongoing project well beyond the life of the EC SynEx Project.

This report, although a final report as far as the EC project is concerned, is really a description of the first phase in establishing a centre of healthcare excellence. New EC Fifth Framework project funding has already been approved to enable new and innovative technology solutions to be added to the work already established in north London.

Rationale for the Components

Health services internationally are grappling with the evolution from paper-based records to comprehensive electronic healthcare record systems.

Present day clinical information systems in hospitals and in general practice are not yet adequate for the challenges of delivering effective and evidence-based healthcare, in which teams of clinicians on different sites are working in partnership and collaboratively with patients [1]. Electronic and paper healthcare records are held in islands of information in independent information systems, each with its own technical characteristics and view of the healthcare domain. The available evidence on good clinical practice, existing as publications and guidelines, is often too generalised to be applied to individual patient groups, and is isolated from the relevant known facts about any particular patient's medical and social background [2]. This evidence is difficult to retrieve at the time and in the location where needed. It has been estimated that 15% of the resources of an acute hospital are currently spent in gathering and processing information, accounting for up to 25% of doctor and nurse time [3]. So there is clearly much scope for improving the efficiency of clinical information management.

The primary purpose of the healthcare record is to support patient care [4], which subsumes many different functions including the documentation of a historical narrative account of the health care given and supporting the communications between healthcare professionals [5].

Healthcare professionals need to document and to review increasing volumes of information, as patients receive increasingly complex care, involving a range of data-intensive examinations, investigations and treatments. The majority of clinical detail in records and in communications, on which such care depends, is still on paper [6, 7]. Clinicians need to share healthcare information with a wider range of professional colleagues on multiple sites. The care of any one patient can require a healthcare professional to review the information held in several distributed clinical systems in a consistent manner [8]. They also need to keep rigorous records to demonstrate their competence [9, 10] in case of future litigation, which is increasing in frequency each year, and in order to justify their use of healthcare resources.

Clinicians and patients also require high quality tools to analyse the data within individual electronic healthcare records [11] in order to:

- observe trends and patterns in the health of a patient;
- enable the use of clinical guidelines and decision support systems as part of evidence-based practice;
- perform clinical audit;
- inform management and commissioning decisions;
- support epidemiology, research and teaching activities.

The need for a clear organisational structure to healthcare records, whether on paper or on a computer, has been discussed on many occasions [12, 13, 14, 15], but it has proved difficult to encourage clinicians to abandon paper records in favour of a fully computerised healthcare record system [16, 17, 18]. There are clear advantages to clinicians themselves of well-organised electronic healthcare records for improving the completeness of the information they elicit from the patient [19], the scope for the subsequent analysis of that data [20], for supporting shared care between clinicians [21, 22], and to demonstrate clinical competence [23, 24]. Improving the ease with

which EHCR applications can be learned and used [25] and developing more efficient means for clinical data entry [26] have both been suggested as solutions. However, the inherent diversity and complexity of medical data [27] and the need to use rich and varied descriptive terms [28, 29, 30] are held by many clinicians to be a fundamental obstacle to adopting more formal recording structures. The lack of an appropriate architecture for healthcare records has been identified as a major impediment to progress in this area [31].

Clinicians and patients need high-quality information systems to connect an individual person or healthcare record to a repository of the knowledge that is relevant to the care provided. This should include relevant background information about the patient (from their healthcare record or records) and evidence-based management protocols that can be directly applied to the care of individuals.

The Health Telematics Research and Development programme of the European Union has recognised many of these health informatics challenges and sought to address them on a large scale through a set of multi-national projects over the past decade [32, 33]. The EHCR server developed by UCL builds on a long pedigree of such European research, principally the Good European Health Record (GEHR) and Synapses projects. Much of the work and experience gained in these projects has informed progress on standards through CEN and now in the International Standards Organisation. A four-part EHCR pre-standard (ENV 13606) was published in June 1999 [34].

From the results of these background projects and standards, a number of requirements are applicable to the information model and the services provided by the federated EHCR and decision support system implemented within the London SynEx demonstrator. These are being documented in full by UCL and will be published later.

The broad areas of requirement are listed in the tables below.

GENERAL FUNCTIONALITY
A generic, open and standards-compliant means for combining healthcare records consistently, simply, comprehensibly and securely, to enable the sharing of data between different information systems in different places.
Generic methods, hardware and software products to enable individual computer systems or devices to exchange healthcare data with each other whether on the same site or accessed by secure telecommunications links.
Generic methods to allow clinical protocols to be implemented as management algorithms applied to underlying patient EHCRs, to generate alert messages or additional EHCR entries.

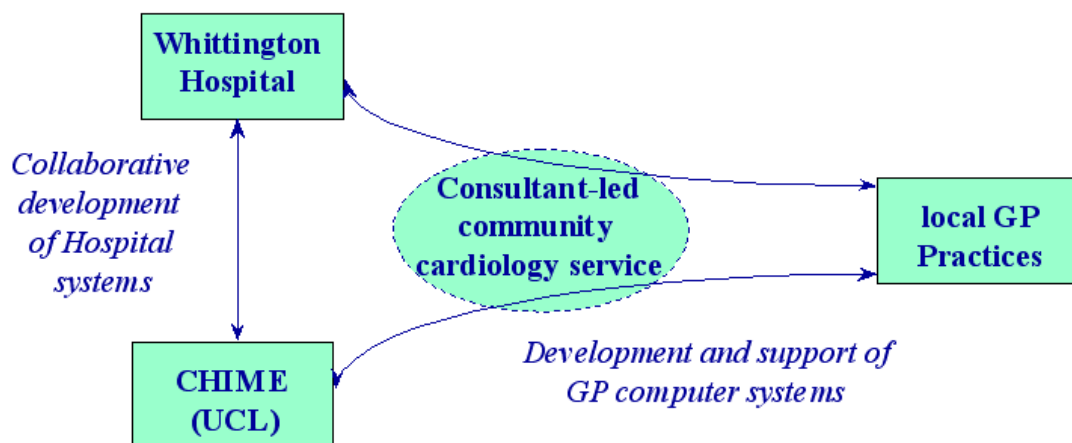| ETHICO-LEGAL ISSUES |
| --- |
| Accountability, attribution |
| Access rights, amendment rights and confidentiality |
| Federation of access rights frameworks |
| Rigorous patient identification |
| Healthcare professional identification |
| Integrity & permanence of patient records |
| Versioning, duplicate persistent storage |

EHCR ARCHITECTURE REQUIREMENTS

Accommodation of all required data types:
- term sets (source organisations and versions)
- free text (natural languages)
- charts, tables, diagrams
- accuracy, normal ranges, instrumentation
- multimedia-media (e.g. images signals, drawings)

Flexibility for dealing with *ad hoc* (loosely structured or unstructured) entries

Preservation of original meaning when data is transferred from its original (feeder system) architecture

Preservation of context within and between record entries:
- original entry groupings
- qualifiers, including uncertainty & negative findings
- internal (labelled) links, references, views
- protocols, references to external knowledge
- minimum grouping of record entries for safe transfer
- template frameworks

Clinical Functionality Requirements

Support of searching, filtering & analysis requests, including:
- the record object request process
- filters to customise overviews of record
- retrieve personal/team/speciality/problem data
- automatically generate alerts, messages, reports
- enable audit

Allow monitoring of patient progress, effect of interventions

Provide rigorous basis for decision support, links to knowledge

TECHNICAL IMPLEMENTATION REQUIREMENTS

Feeder systems

Feeder sign up

Managing the clinical database federation

Diversity of feeder systems

Maintaining mapping schemata

Optimising object reuse

Local data repositories: synchronisation with the principal EHCR database

Decommissioning of feeders

Persons and Clients

User registration (single sign on)

Patient ID reconciliation

Client application registration

Adaptation of clients

Security

Secure communications channels and encryption tools

Audit trails

Access and amendment records

Alert and disease management message register

3rd party disclosure register

Authorisation and authentication

Access control management

Time out management

Other

Interfaces to other components

Performance

Exception handling systems

Backup systems

These requirements have formed the basis of the UCL SynEx components, which build in particular on the R&D results of GEHR, EHCR-SupA, GAMES and Synapses.

The Demonstrator Site Partners

The London SynEx Demonstrator comprises a set of primary and secondary care sites in north London working in partnership with University College London. The eventual shape of the demonstrator site will comprise the following healthcare settings.

➢ The Department of Cardiovascular Medicine at the Whittington hospital, integrating:
  ⬧ anticoagulant clinical management system
  ⬧ coronary artery database and display application
  ⬧ atrial fibrillation specialist systems
  ⬧ cardiology investigation and monitoring devices (Marquette/General Electric)
➢ 2-4 community-based consultant cardiology clinics
  ⬧ communicating within the federation through secure telecommunications links
➢ Several GP practices in north London
  ⬧ integrating data from GP-CARE (RFA 4Plus accredited GP computer system)
➢ A new Centre for Chronic Illness to be opened next year, offering integrated ambulatory patient care, providing:
  ⬧ comprehensive EHCR system
  ⬧ links to electronic guidelines and educational libraries

Figure 1: Partner sites in the London SynEx Demonstrator

The initial pilot site for the UCL SynEx components is the Whittington Hospital anticoagulation clinic. This demonstration environment is being used to provide "proof of concept" and to validate initial user acceptance before rolling out access to the services to the wider clinical community.

The Whittington Hospital

The Whittington is a community based teaching hospital in north London, serving a busy and cosmopolitan part of the capital city. It was once the largest in Europe. Through reciprocal links with University College Hospital it is able to offer a comprehensive and expanding range of patient care services. The hospital has a close relationship with University College London for medical student and postgraduate education, and with Middlesex University for the training of nurses and other health professionals. The consultants have close working relationships with local GPs, and the hospital provides support for many GP educational activities.

The Department of Cardiovascular Medicine has three consultants and seven junior medical staff supported by several cardiac technicians and specialist cardiac nurses trained in anticoagulation and in resuscitation. The team provides care for around 1,800 emergency inpatient admissions, almost 7,000 outpatients and co-ordinates around 20,000 cardiac investigations per annum.

The Department's close working relationship with the Middlesex Hospital (within the UCL Hospitals group) for tertiary referrals embraces cardiovascular medicine and cardiovascular surgery, and incorporates peripheral vascular as well as cerebrovascular disease. The recent appointment of a Senior Lecturer in Community Cardiology will facilitate the development of a seamless cardiology service from the patient's home through primary care to secondary care and to tertiary care.

North London General Practice

The GP-CARE project is a collaboration between the Centre for Health Informatics and Multi-professional Education (CHIME) and a set of general practices in north and east London. The project focus is on the design, implementation, support, maintenance and evaluation of high-quality GP computer systems for the improvement of patient care and for practice management. The team has written a completely new cutting-edge GP system. This has been accredited to the latest national standard.

Community Pharmacy

A selected group of north London pharmacy sites will also pilot the use of the SynEx components, accessing the services via a secure internet connection within NHSnet. Commercial sponsorship arrangements for this are still in negotiation.

The Overall Vision for the London Demonstration Site

The long-term aim of the London SynEx implementation is to provide access to the patient information collected around the Whittington hospital and particularly within the cardiology department so that authorised clinical personnel can get access to it from wherever they are working. Staff often need to roam within the hospital, being at times within their principal department and at other times doing clinics or ward rounds. GPs and community teams will need to access hospital-based records, and to take advantage of shared care records and protocols.

The main objectives for the SynEx clinical data integration in London are to:

- provide a seamless service across primary, secondary and tertiary care and for self-management;
- interface seamlessly with current systems within and without the hospital setting;
- support the continuing development of systems such as the anticoagulant control advisory system;
- integrate the coronary artery database and display;
- give multimedia access to electronic medical libraries and repositories of best practice;
- support the development of educational and behavioural change material;
- support clinic and department management;
- receive and process data from cardiology equipment.

The SynEx London demonstration incorporates:

(a) the federation of legacy patient records from the primary and secondary care sites described above, to provide an integrated view of an individual patient's healthcare history and to support clinical audit across population sub-groups;

(b) the creation and communication of new record entries through common template structures and through the use of protocols;

(c) the use of an expert advisory components utilising research-validated algorithms and reference data-sets to calculate treatment and monitoring information for managing each patient's condition;

(d) the integration of the ICRF's PRO*forma* guideline authoring and delivery software to deliver cardiology guidelines in the Department of Cardiovascular Medicine at the Whittington Hospital and in local community based settings.

The clinical content of the demonstrator prototypes will focus on the cardiology and general patient summary information in the various Whittington Hospital cardiovascular departmental systems described below (legacy data).

The north London demonstrator vision is to deliver the seamless shared care of patients with cardiovascular illness, in a managed care environment. Patients requiring anticoagulation will have their therapy commenced in the hospital outpatient clinic, informed by background information from their GPs record, guided by electronic protocols and decision support systems. Once stabilised, their care will be

transferred into the community and managed under the same protocols and their GP having appropriate access to the hospital records. The care of all patients under anti-coagulation will be subject to clinical and management audits, through interrogation of their federated health care records. Similar managed care scenarios will be developed over time for patients suffering from angina, ischaemic heart disease and cerebrovascular disease, with a particular focus on protocol-based therapy guidance.

There are several other plans to introduce new features for information delivery and patient management, meeting new requirements for transoesophageal echocardiography, vascular and rehabilitation work, and epidemiology insights.

Recently the Whittington has been also been pioneering the use of "Integrated Care Pathways," presently paper based mechanisms for automating the path of the patient across departments in the hospital as activities are performed and as a disease changes in nature. There is undoubted utility in making these forms computer based, not least because they are currently not particularly readable.

It is well known that patients can acquire considerable expertise in managing their own health if they are given useful material with which to educate themselves. There is a compelling argument for the benefits of such "passive" decision support (advice that is not patient specific). In the longer term SynEx components will relate parts of a patient's record to relevant advice topics, delivered through a combination of web-based stations and touch-screen kiosks.

Existing (Legacy Systems) Environment

There is a need to access the proprietary hardware and software systems that together make up the cardiac care equipment in use in the department. Some of these systems are used in the Middlesex Hospital, which is physically distant from the main Whittington base. At present the cardiology department has several heterogeneous computer systems with relatively limited inter-operability:

1. an outpatient system collecting information which includes outpatient diagnoses;

2. an inpatient system which is used to a limited extent for logistic rather than technical reasons;

3. a test system for echocardiography, 24-hour tapes and exercise tests;

4. an anticoagulant advisory (decision support) and patient management system;

5. a coronary artery disease system (incorporating some decision support features) which is presently being brought into clinical use and will be evaluated in that context;

6. a Marquette/General Electric/Helliger system which embraces:
   • an expert and advisory system for ECG interpretation;
   • an ICU system with six monitors;
   • a CCU system with six monitors and a central station;
   • an exercise test system;
   • a 24-hour tape system;
   • ECG machines.

All ECGs recorded in the Hospital, whether they be inpatients, outpatients or in A & E Department are recorded electronically and stored. Although access to previous individual records is therefore available, the access is laborious and not instantaneous. The departmental workstations are connected via a Local Area Network to the Whittington HIS. The anticoagulant system is used within the hospital outpatient department and (via laptop computers) in community clinics.
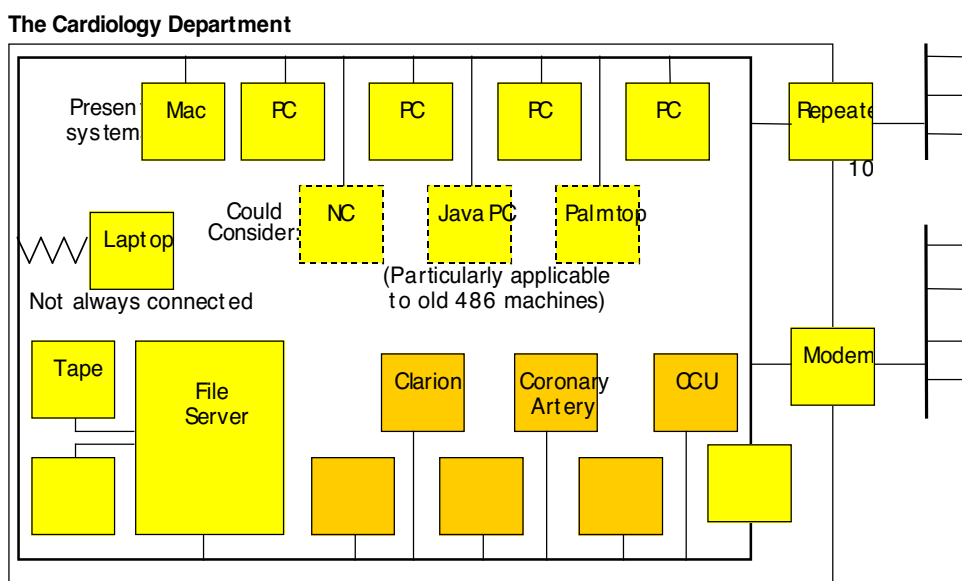
Figure 2: Overview of the Present Cardiology Departmental IT Systems

The Department of Cardiovascular Medicine intends to continue its present arrangement with Windows^{TM} client workstations communicating on a LAN running Windows^{TM} NT and Novell® Netware^{TM}. The GP-CARE practices use a mixture of Novell and NT networks to connect Windows^{TM} workstations. All sites will also be capable of communication within the NHS-wide network via approved security mediators, and will be capable of exchanging EDIFACT messages with other parts of the NHS. The north London demonstrator will eventually federate the patient records from all of these legacy systems.

The SynEx London Demonstration: Distributed Anticoagulation Services

The work of implementing new SynEx components, customising them for deployment in the north London cardiac care setting, putting them into operational clinical use and carrying out appropriate evaluations is expected to continue for some years, well beyond the life-time of the present EC funded project. The commitment of the local provider organisations to the SynEx approach is therefore a significant one. The first stage of this demonstration is in the field of anticoagulation, which has been completed and made fully operational during the EC funded project lifetime.

The Whittington Hospital cardiology department already undertakes world-leading research in fields such as anticoagulation. However, the legacy anticoagulant advisory and management system still has relatively restricted access. The results of the routine INR blood test are entered on paper records in the laboratory and brought physically by the patient to the anticoagulant system where they are entered. The anticoagulant advisory system, written in Access Basic, is not available on-line to the wards. The intent of the SynEx demonstrator is to facilitate on-line access to a new electronic advisory management system and to minimise the need for both manual and electronic entry.

The demonstrator is utilising the following UCL components:

- **Object Dictionary Client and services** (as a means of facilitating feeder system sign-up and of navigating a federation environment)

- **Federated Healthcare Record services** (as a scalable run-time environment supporting distributed access to record components from new and legacy feeder systems)

- **Persons Look-up services** (for patient demographic information and staff identifiers)

- **Expert Advisory (Decision Support) services** (to calculate the patient's next treatment regimen and next monitoring interval)

The key <u>deployment</u> achievements to date include:

- designing and developing the Object Dictionary Client (ODC), and authoring the anticoagulant object dictionary schema with mappings to feeder system schema
- implementing the primary record component database: porting of feeder system data from the legacy application to ObjectStore using the full standards-conformant SynOM;
- generating end-user clients through web-servlet-based extraction methods that provide the rich functionality of a full clinical application;
- implementing the Persons Dictionary using Novell NDS, and populating it with the 39,000 NHS GPs, the north London hospital consultants, and the patients presently registered with in the legacy anticoagulant system.

These components are summarised below and documented fully in other reports. The implementation work is almost complete, and the clinical staff are now beta testing the new system with a view to its live use later this month.

A formal evaluation framework is being developed to enable the components to be properly assessed prior to their wider distributed access. The metrics wll include:

- the rigor of the information models underpinning the FHCR and Object Dictionary
- the design of the ODC, and its representation of the anticoagulation schema
- the faithfulness of the feeder system sign-up
- the robustness and performance of the FHCR services
- the accuracy of person identification
- the clinical utility of a distributed clinical application

A formal set of requirements are being collated from work commenced in GEHR and refined during the Synapses project. These will be used to develop a set of test criteria.

Anticoagulation Client Screens

This successor application has drawn on an existing Visual Basic application, and initially provides a set of HTML web clients. The overall application includes forms to deal with requests for and the display of existing data, and also with data entry. Java applets will in the future be used as necessary to provide end-user graphic objects not supported by the core HTML standard.

Some example client screens are shown below, captured from a web browser as they are in practice being accessed by the clinical team. Each screen has a brief caption to outline its role. A fuller set of the screens, with brief explanatory notes, is given in Annexe A of this report (provided as a separate document).
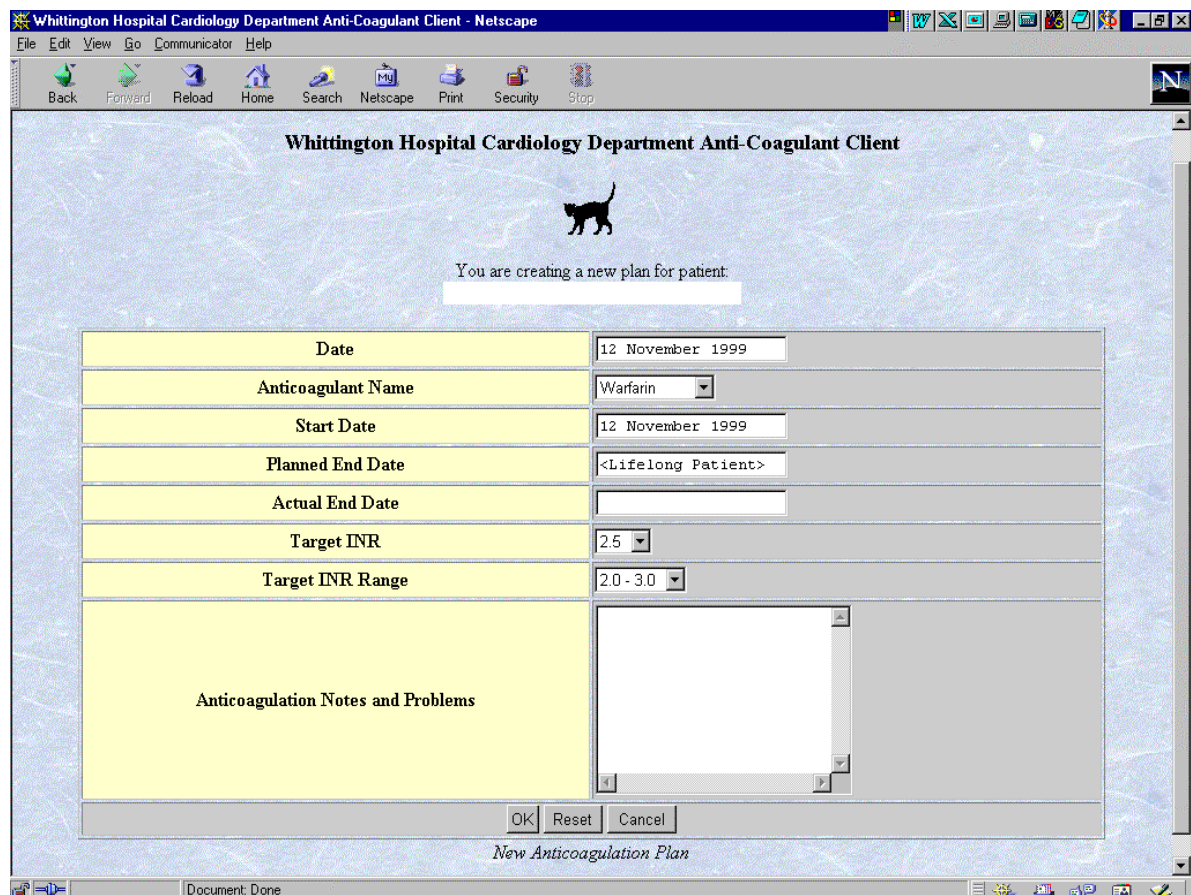

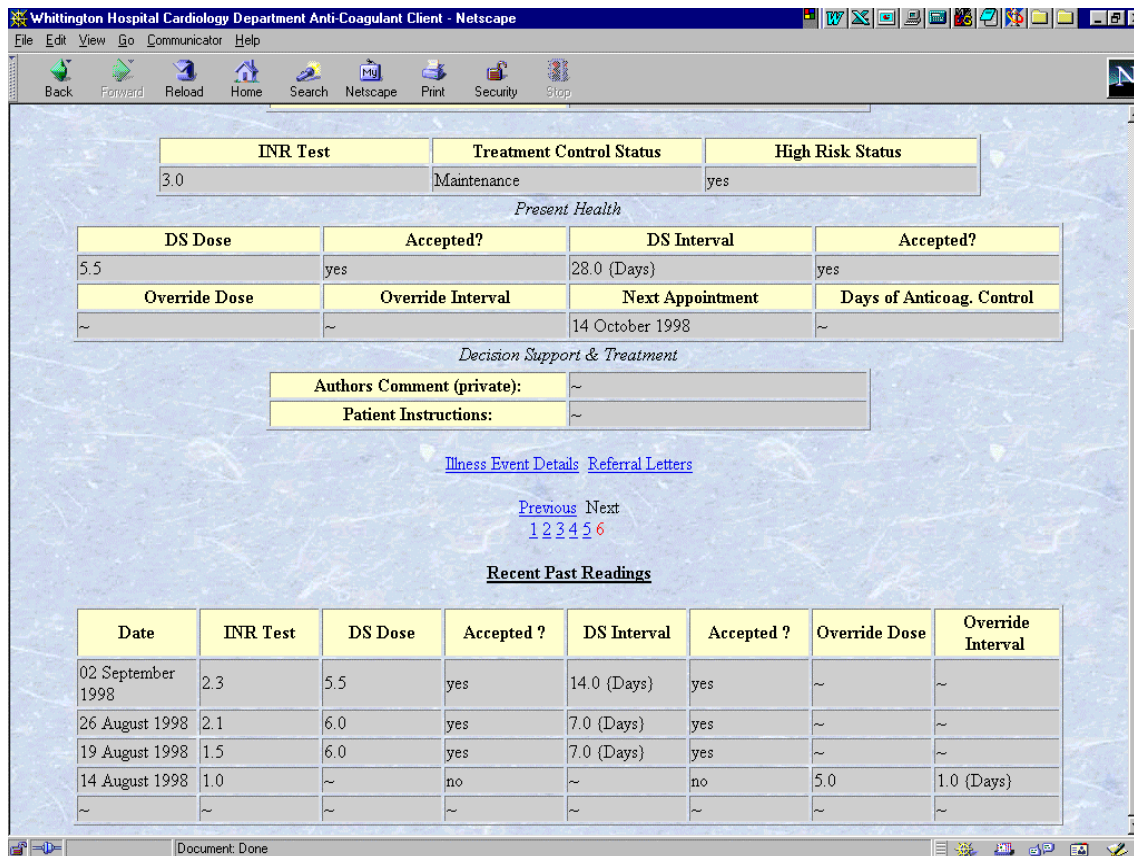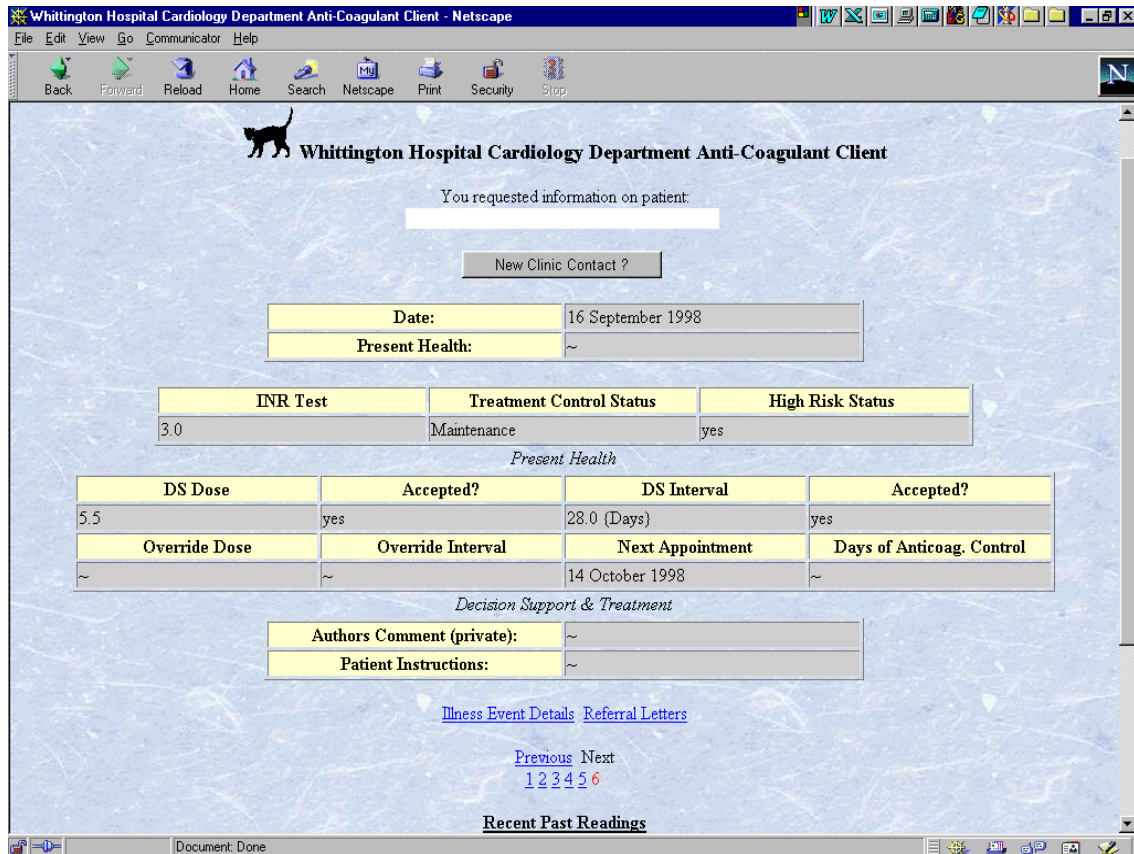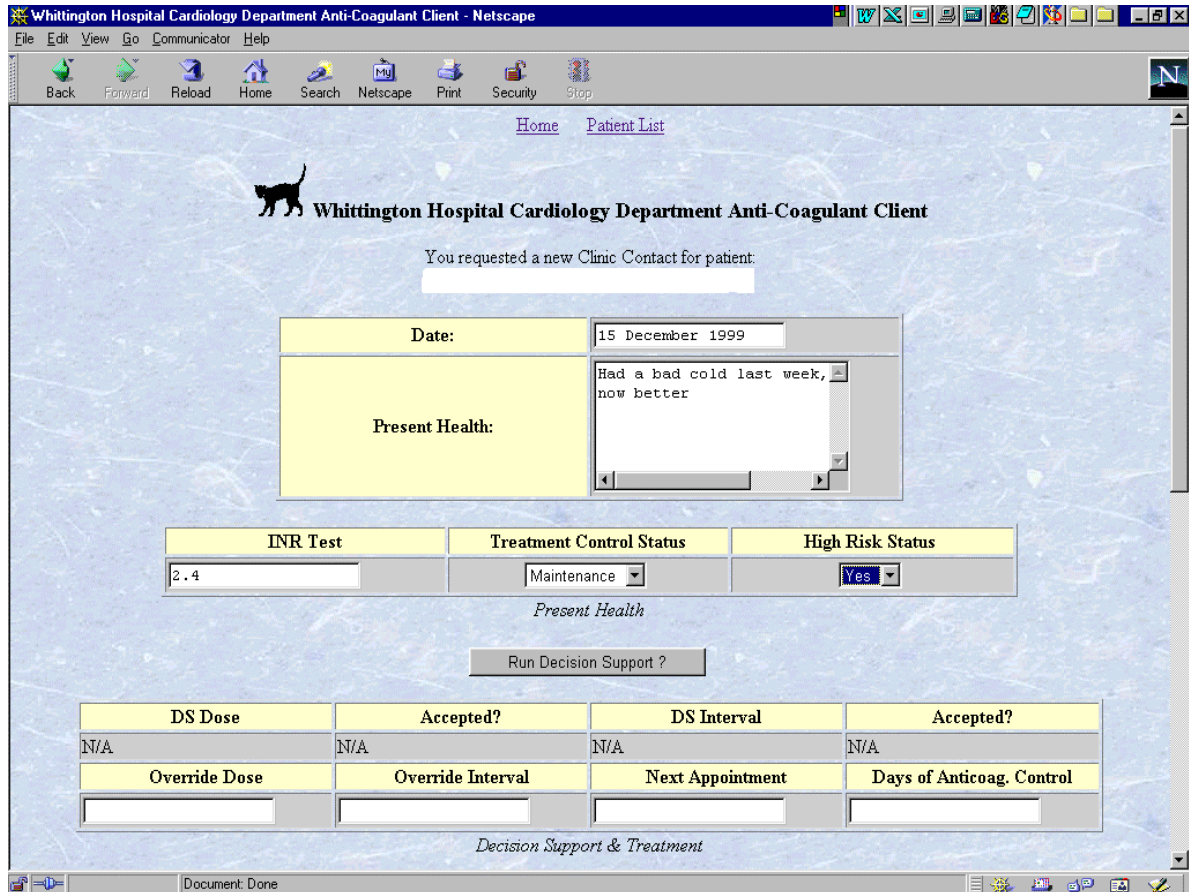**Figure 3: Anticoagulant Client - authoring a new treatment plan**

Figure 4: Anticoagulant Client - viewing a clinic contact

(top and bottom half of the web page)

Figure 5: Anticoagulant Client - entering a new clinic contact

(prior to running the decision support)

Technical Summary of the Demonstrator

Overview

The SynEx prototype systems developed and evaluated in north London demonstrate:

- the comprehensive use of EHCRs for the management of cardiovascular care through ENV 12265 conformant and ENV 13606 compatible record systems;
- the federation of existing legacy systems through a Synapses object dictionary containing local and national standard data-sets mapped to the schema of each feeder system;
- the shared adoption of clinical management protocols through the use of a PRO*forma* -based protocol engine with SynEx-compliant interfaces to the Synapses Object Dictionary.

The development of client-server applications for large enterprise-wide intranets or extranets has become possible because of the recent availability of solutions and products. To ensure an optimal opportunity to market the SynEx components within a multi-vendor community of EHCR systems developers, these functions will be offered through a set of component-based services integrated within a platform-independent and standards-based distributed middleware environment.

The set of middleware components intended for this demonstrator is shown in Figure 6 below. In functional terms, they will provide:

- a comprehensive, multi-media, multi-professional electronic healthcare record system
- the secure and seamless federation of disparate record systems within and between sites
- clinically-oriented tools to facilitate the creation, navigation and analysis of patient records
- integrated terminology, protocol and decision support services
- links to a range of clinical knowledge and educational databases

Each of the UCL components or services is described in more detail in the rest of this section. All of the main components are written in Java™, and exist within a middleware environment managed through directory services and JINI™, an open standard service-integration technology. This will allow the ongoing development of flexible and portable applications with high-level graphical user interfaces to be made where such applications can inter-operate across diverse architectures and infrastructures.

In Figure 6:

- components shown as yellow objects provide or support core record services
- components shown as blue objects service or represent various aspects of clinical knowledge
- components shown as green objects are non-middleware client applications that will use the core services
- red font has been used to show existing products or advanced prototypes of other consortium members, which may be deployed in live or test scenarios within the demonstrator

The live demonstrator will include at minimum the CHIME components in blue font plus PRO*forma,* and in the future GALEN services. The other components (DHE, I4C and NUREC) will be demonstrated in a non-clinical setting (i.e. with test data) as proofs of concept, although this may only be possible at the end of the project. The London site will seek where possible to provide a means by which these other products can be demonstrated to potential purchasers.
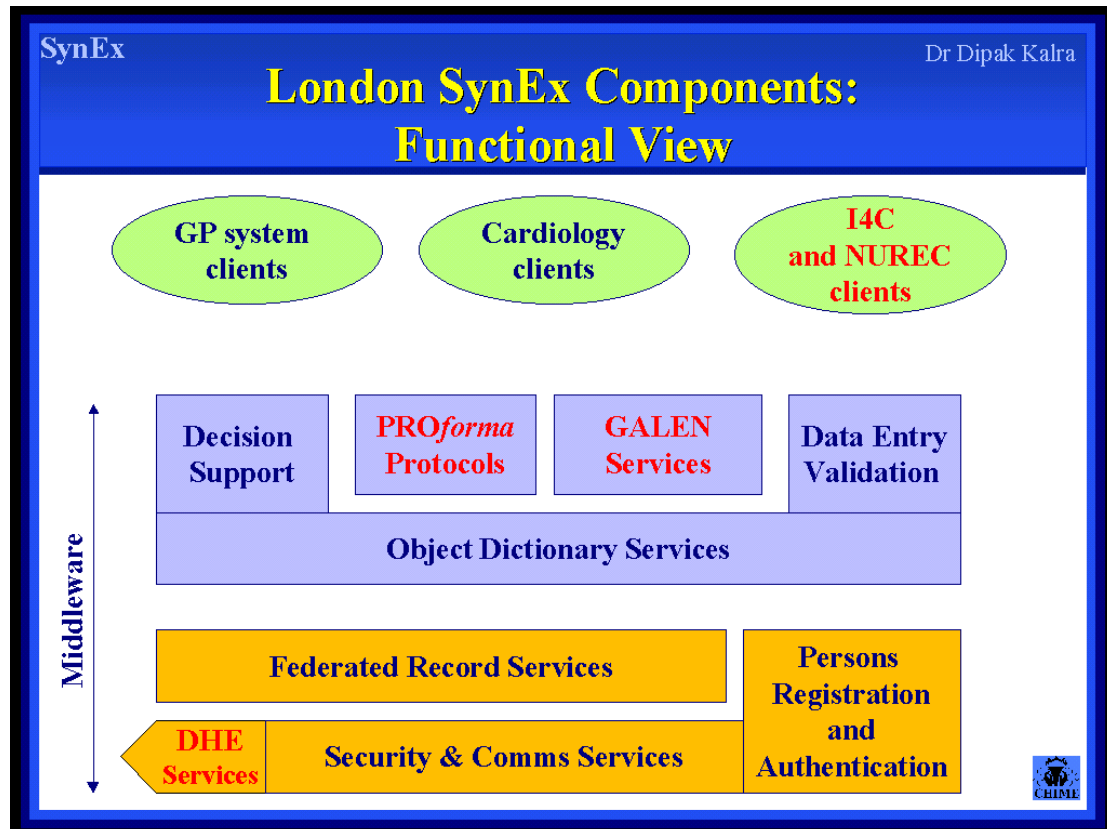
Figure 6: Overview of the Middleware Components Intended for the London Demonstrator Site

The information models and interface specifications for each UCL component will be offered into the public domain at the end of the project, through the public deliverables of the SynEx project and through academic publications.

The diagram below shows the schematic relationship of the various components and sub-components providing distributed access to federated healthcare record services.

Middleware Computing Environment

The heart of the London SynEx demonstrator is a set of directory services accessed through the Java Naming and Directory Interface (JNDI) and a distributed set of record and other expert knowledge services delivered via JINI^{TM}, providing the run-time access to:

- the Synapses Object Dictionary
- a set of legacy and newly developed data feeder systems
- an EHCR Object Repository
- a dictionary of persons and devices
- a dictionary of access permissions
- access to other knowledge services (e.g. decision support, protocols)

These are shown schematically in Figure 7 below.
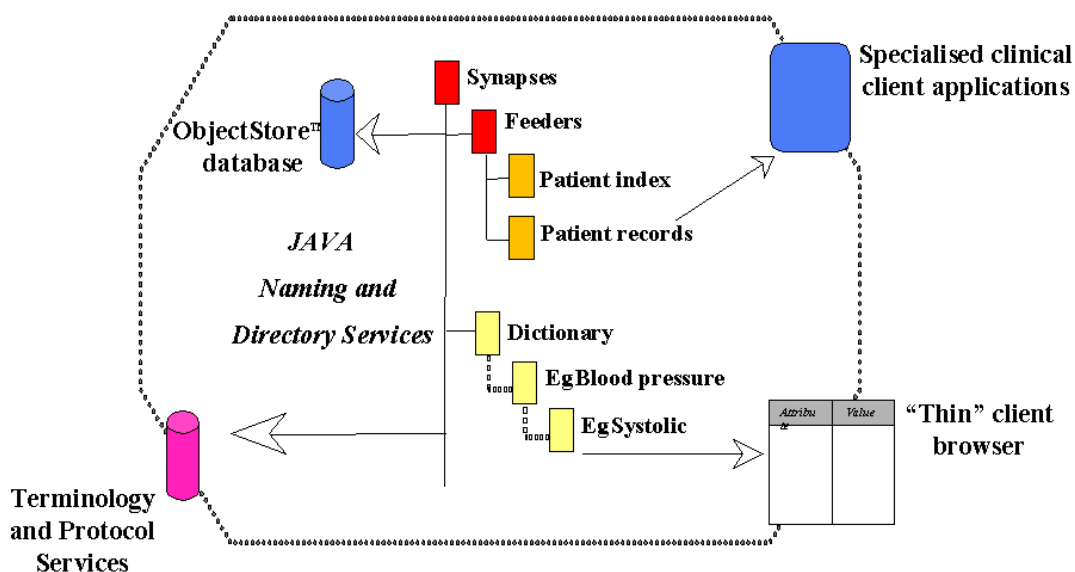
Figure 7: Data services delivered through Directory Services and JINI

Directory Service Approach

The federated access to distributed clinical databases is managed through a set of directory services accessed via the Java Naming and Directory Interface (JNDI). This environment provides the run-time access to the record objects defined within the Object Dictionary Client (ODC), drawn from legacy and newly developed data feeder systems.

The whole computer industry in general is investigating the problem of locating and federating distributed data. In particular, Novell, Netscape, Sun and Microsoft have considered rival technologies as an answer to the problem of locating information about people and things. Netscape and Microsoft have favoured a technology called LDAP (Lightweight Directory Access Protocol) for accessing directory-oriented information. Novell have provided an LDAP interface to their own NDS directory system. Sun has created a technology for Java called JNDI (Java Naming and Directory Interface) that can access any directory-oriented service including LDAP. Directory Services are therefore now becoming the industry preferred method of locating objects within containment hierarchies.

It is important to note that the LDAP and JNDI are not aimed at database access, but at the arena of enterprise naming schemes. The directory service is accessed in a completely uniform way, independently of the type of storage (a file system, a database, or anything else can be represented as an object) and the data can be accessed irrespective of the database's preferred access mechanism (SQL, OQL, etc.). A particular client may use SQL to access and update a particular feeder system but JNDI gives the power to extract data however it is stored. It also provides a way of getting the database schema should it be required.

Many federated object sources can be attached to a hierarchy, and can return objects and attributes from a lookup. Feeder system "signup" *is* the process of attaching objects to the directory. Any object source on the network can attach to the directory; Synapses federation therefore follows from the use of the service. Any client that can see the directory automatically has access to the whole Synapses Object Dictionary and patient record databases (within appropriate security frameworks).

JNDI provides a uniform access to the set of SynEx *data* component services. These and other SynEx services (such as decision support and terminology)) have been successfully integrated and presented for client access through JINI (a new technology delivering seamless and unsupervised access to services, see below).

The directory service component set comprises a specification for the configuration and deployment of these emerging industry standards, together with tools to facilitate the process of feeder system sign-up. The demonstrator shows how patients, clinicians and the anti-coagulant EHCR record objects can be accessed in this way.
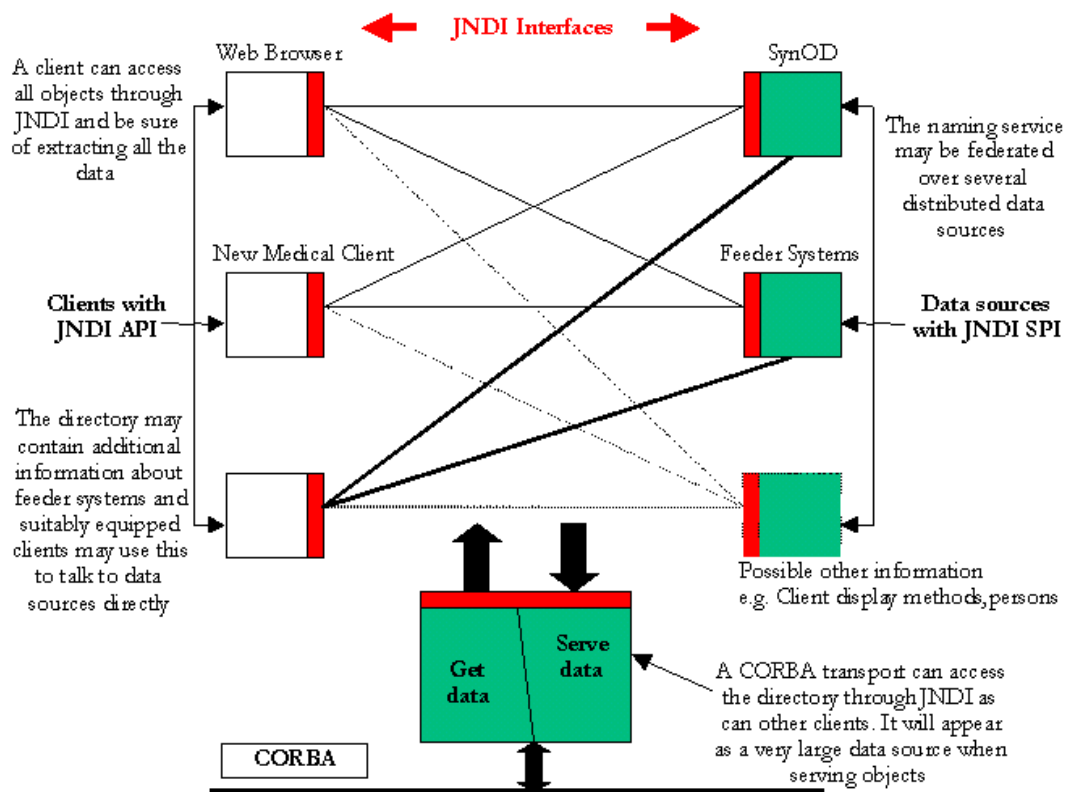
Figure 8: JNDI Interfaces to Clients, Data Sources and to a CORBA Transport Layer
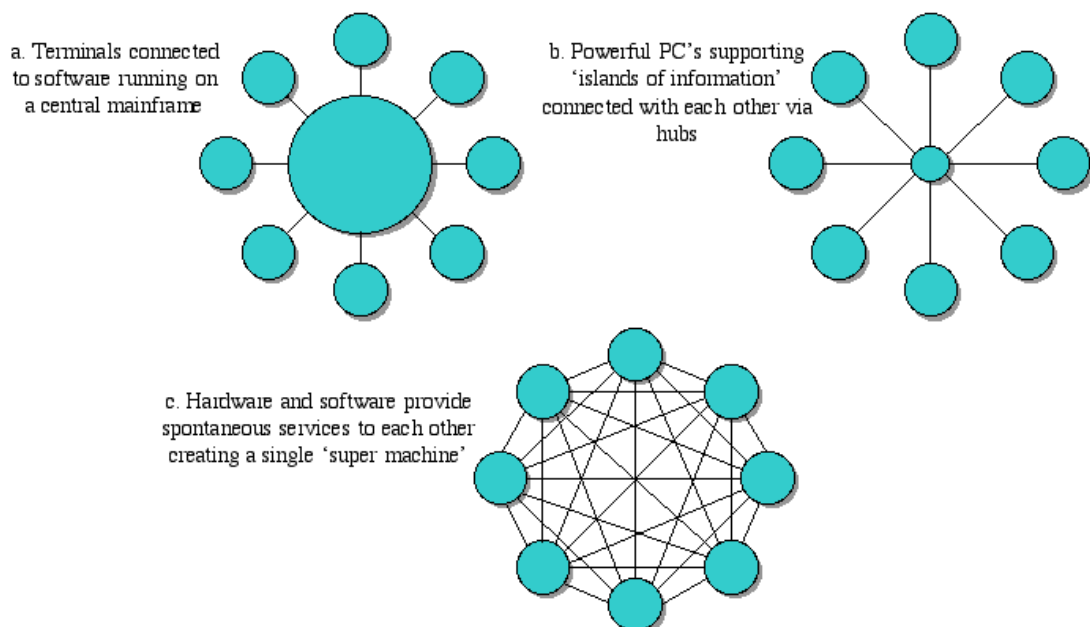
JINI

JINI, pronounced 'genie', is the latest API from Sun and promises to be nothing short of spectacular. The idea behind it is relatively simple and in fact several companies are known to have been working on something similar. But they have all had to wrestle with the architecture independence issue. Either one must assume the ubiquity of ones own technology, which other manufacturers are reluctant to do, or some way must be found to abstract every architecture.

JINI provides a mechanism by which any item of hardware or software can make itself available to every other item on a network without any intervention from a human network manager. The idea is a derivative of the networked PC model of computing where every aspect of the PC is either able to join the network in its own right or be proxied by something else which is. The first and most obvious conclusion is that the need for device drivers is far less because everything has a standard way of making itself available. This should reduce the problems associated with device driver failure such as non-functioning printers and so forth.

The second conclusion however, and more pervasive, is that even 'non-computer' machinery such as heart monitors etc., can be added as services on a network and can demand services from other machines. The monitor might demand access to patient record software to add data to it while a cardiologist attaches the output to a visual display unit on his desk so that he can keep watch on the patient while working elsewhere.

For interest, consider the evolution of computing as shown in figure 2. JINI really is a new paradigm for the provision of services in a workgroup. We can envisage that a class file is a unit of development for a software engineer, a JavaBean is a unit of assembly for a system provider, and a JINI service is a unit of usage, a component that users will be comfortable with.

a. Terminals connected to software running on a central mainframe

b. Powerful PC's supporting 'islands of information' connected with each other via hubs

c. Hardware and software provide spontaneous services to each other creating a single 'super machine'

Figure 9: The three evolutionary stages of computing

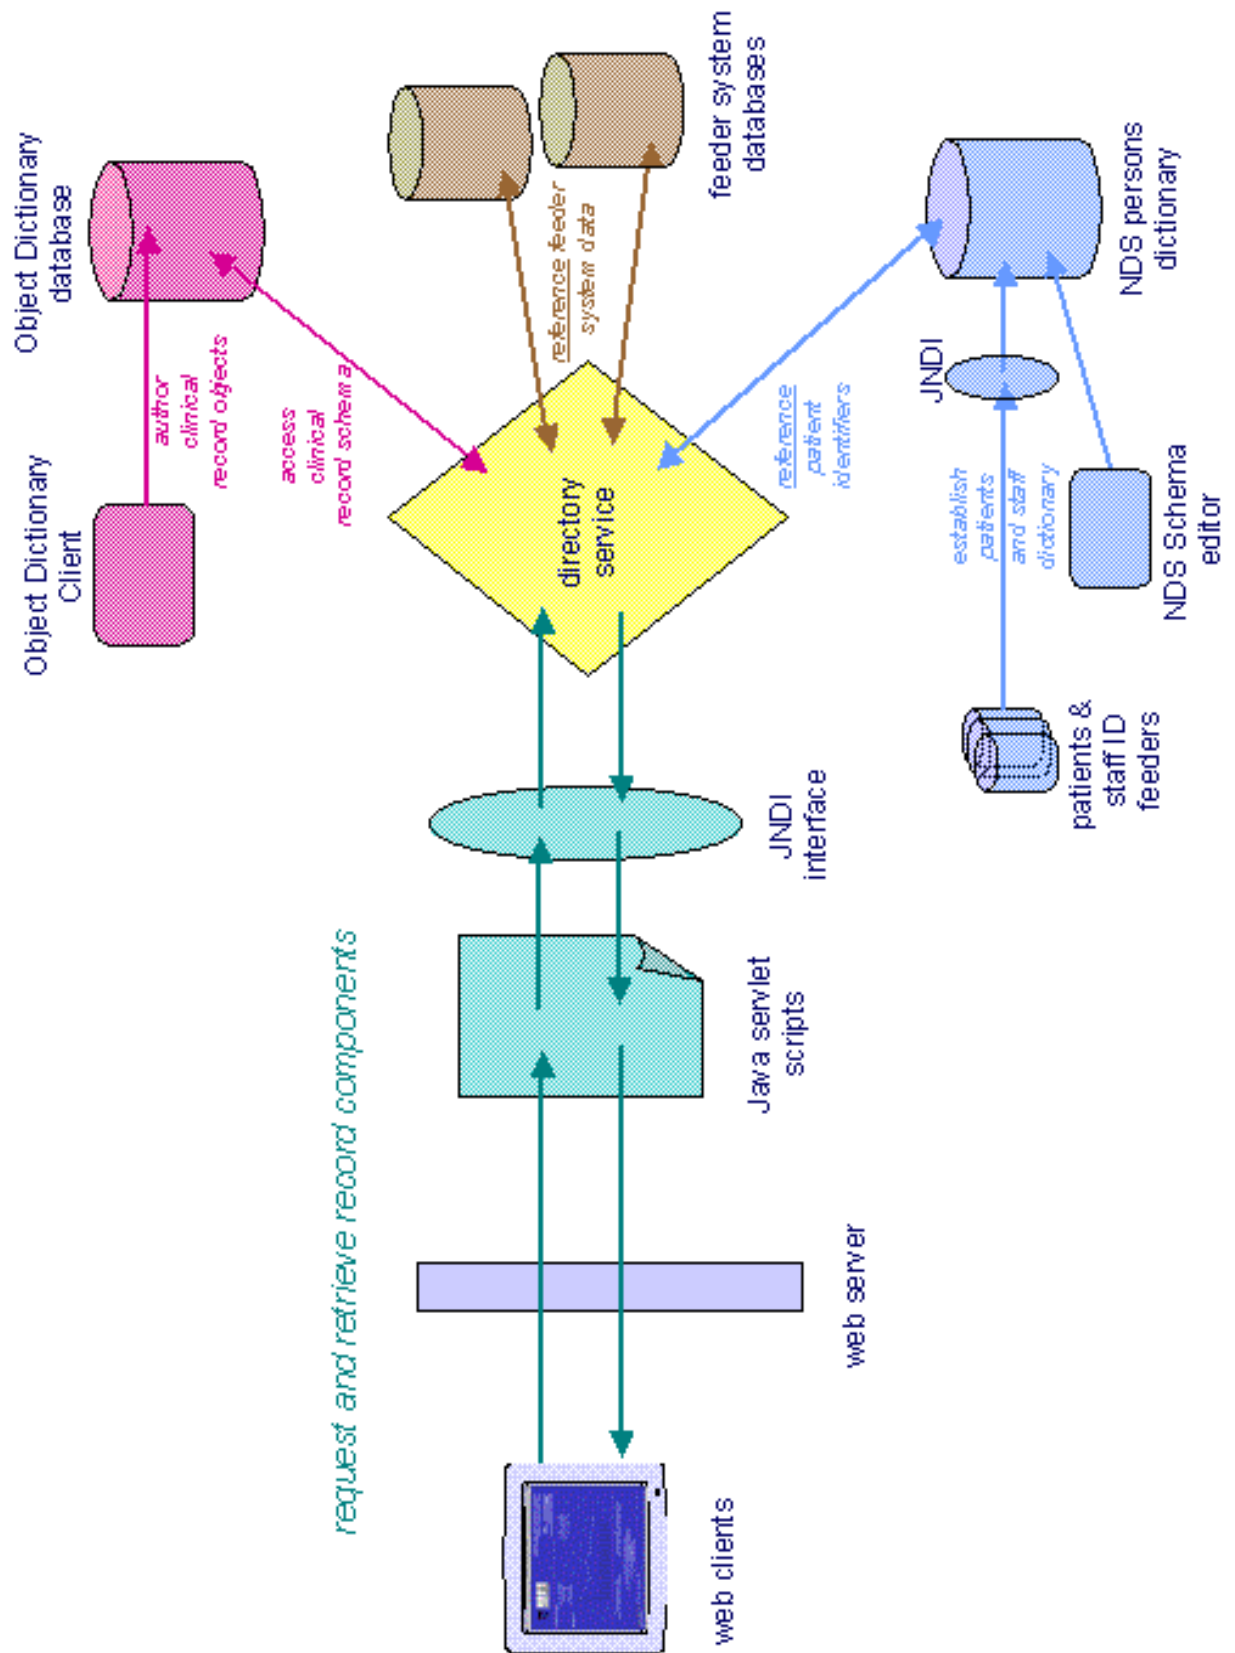## *Federated Record Service Component Set*

Figure 10: Engineering Overview of the FHCR Services

The federated healthcare record is a high-level abstract model, enabling the federation of records from a diversity of feeder systems. In essence, it is the generic schema to which all federation feeder systems are mapped.

The federated healthcare record architecture adopted for the UCL components and for the London demonstrator slightly extends and refines the model developed by the Synapses project (the SynOM), informed by the work of EHCR-SupA[1]. It is based on a rigorous object model formalism; it incorporates the representation of content types and placeholders for ongoing work on the integration of medical knowledge (GALEN) and protocol (PRO*forma*) services. It also enables the generation of record extracts and messages conformant to the latest CEN/TC251 standard for Electronic Healthcare Record Communication: ENV 13606.

The FHCR architecture has been implemented as a set of Java classes (and capable of conversion to an XML DTD) that provides a reference model for:

- • the object dictionary (see below)
- • feeder system mapping
- • the EHCR object repository
- • client server communications

EHCR database

As well as accessing distributed feeder systems, the UCL FHCR services incorporate a principal EHCR database that can be used as a local cache and provides a robust repository for data originating from feeder systems that are to be decommissioned. This object oriented stores record components in a form native to the federation architecture.

For functional and performance reasons, Object Store (from Object Design Inc.) has been chosen as the core database of the server environment. In the demonstrator site, the existing anti-coagulant databases were transferred to this Record Object Repository as Java objects. An example set of tools to facilitate the decommissioning of feeder systems and/or the presentation of their data have been developed.

Object Store (from Object Design Inc.) has been chosen for this as it provides:

- compatibility: an object-oriented programmatic interface is provided through Java. Java objects can be stored directly without translation.
- scalability: caching and replication is utilised is order to meet the requirement for large numbers of simultaneous object requests.
- robustness: features such as automatic recovery, on-line backup, roll-forward and failover enable continuous operation.
- tools: schema can be modified, databases populated with sample data and examined through a range of command-line and graphical user interface style tools.

Object Store has been implemented by installing it both on Solaris 2.5.1 and Windows NT 4.0 SP3. Within the framework clients of the EHCR database communicate with a wrapper-class to ObjectStore, insulating them from its exact operation. This wrapper-class is itself an RMI (remote method invocation) listener

---

[1] EHCR-SupA is a Fourth Framework support action that has taken forward the work of GEHR and other EHCR architecture implementation feedback

node executing continuously, so that clients can run in their own processes and/or remotely. Instances of any classes that are stored in ObjectStore need to have been 'post-processed'. This procedure adds and inserts methods to facilitate timeliness and memory-management so that data in the object graph is only retrieved from persistent storage when the corresponding 'get' method is called. The data requirements could be achieved through the creation of just two of these classes.

Web Servlet Approach

A set of web servlet scripts has been written, using Java, to extract single or multiple instances of patient record objects from ObjectStore. The servlets map the output object attributes to cells within html tables. This provides a means to verify the information content of each object, and provides a simple record display, but is not intended to provide a clinically suitable client interface.

Java servlets have become increasingly popular as they enable the functionality of the web server to be extended for the dynamic creation of content. Being written in Java they are secure, cross-platform, re-usable and offer good performance through the convenient use of threads. Engines for servlets exist either as integral to the server or as pluggable modules. Servlets were an obvious choice here with so much already being implemented in Java. The servlets here talk to the database and dynamically generate the HTML that is sent to the web browsers of the system users. Likewise they handle any requests e.g. to add a patient or to change their details, sent by users. Servlets can be written by anyone with a moderate understanding of Java, HTML and the HTTP protocol. Java Server Pages (JSP) provide a higher-level script-based technique for developers, which utilise servlets as their underlying technology.

Anticoagulant Specific Components

Some middleware components have been authored specifically for use in the management of anticoagulation therapy. The existing decision support methodology (i.e. the algorithm and tables for warfarin control) has been re-engineered using Java. This service is now provided through specific agents called from a dedicated client and returning data to this client. Their deployment in other settings will be openly reviewed, but these components are presently only being actively promoted for potential commercialisation as part of the anticoagulant application. Future work on their generalisation will be considered later.

Object Dictionary Service Component Set

The Object Dictionary component provides the formalism by which the specific clinical data sets and aggregates normally found in healthcare records and in contemporary feeder systems can be defined. Any such dictionary entry utilises the SynOM classes as basic building blocks, extending the classes to generate specific clinical hierarchies that can be directly mapped to feeder system data schemata and can be the target of a client request.

The Object Dictionary defines the domain specific classes used in actual healthcare records. It identifies the object sets to be retrieved from federated feeder systems in response to a client request, and incorporates references or access methods to the underlying feeder system EHCR data.

This component provides the tools to allow end-users to author entries in the Object Dictionary for the purposes of mapping "user objects" to the underlying feeder system schemata.

It is the intention of UCL to develop a means by which locally or nationally approved object definitions can be published and shared as a reference library of clinical terms and data sets. This is intended to be through a public and/or restricted web site, with XML as the preferred format.

The Object Dictionary Client (ODC) component has been written entirely using Java Foundation classes and Swing, allowing true cross-platform deployment. It utilises an object database PSE Pro, from Object Design Inc., which is also a Java application and is similarly capable of installation on any platform that supports a Java Virtual Machine. The licence for PSE Pro permits the distribution of run-time versions alongside the Object Dictionary application, removing the need to purchase any additional third-party software. PSE Pro was chosen over Object Store as this allows for multiple users to become involved in authoring or viewing the Object Dictionary, potentially using stand-alone equipment.

The client provides a hierarchical (tree) display and a data entry form, with an evolving set of search and management functions. Objects can be rendered obsolete or formally revised using this tool. The ODC permits the structure of the record object definitions to be captured in a way that the user originally intended for maximum performance and flexibility.

The core features of the ODC are listed in the table below.

ODC Class Hierarchy
ODC Object Properties
Creating New Object Entries
Cardinality on Instantiation
Validation Criteria
Data Retrieval Methods
Copying and Pasting Objects in the Hierarchy
Publicising Objects
Deleting an Object
Marking an Object Obsolete
Revising an Object Definition
Reviewing the Version History
Tracking Objects with Multiple Parents
XML Export of the Database
Saving the Database

Table 1: List of Features within the Object Dictionary Client Tool (ODC)

The authoring tool incorporates an XML generator and will shortly also include an XML parser to support the interchange of definitions between installations or with other sources such as public domain libraries. It also incorporates a replication function to allow for the synchronisation of object dictionaries within a distributed environment, and for the separation of prototype and definitive entries.

The next stage in the development of this tool is the incorporation of Object-Feeder system mapping functions and run-time methods for retrieval.
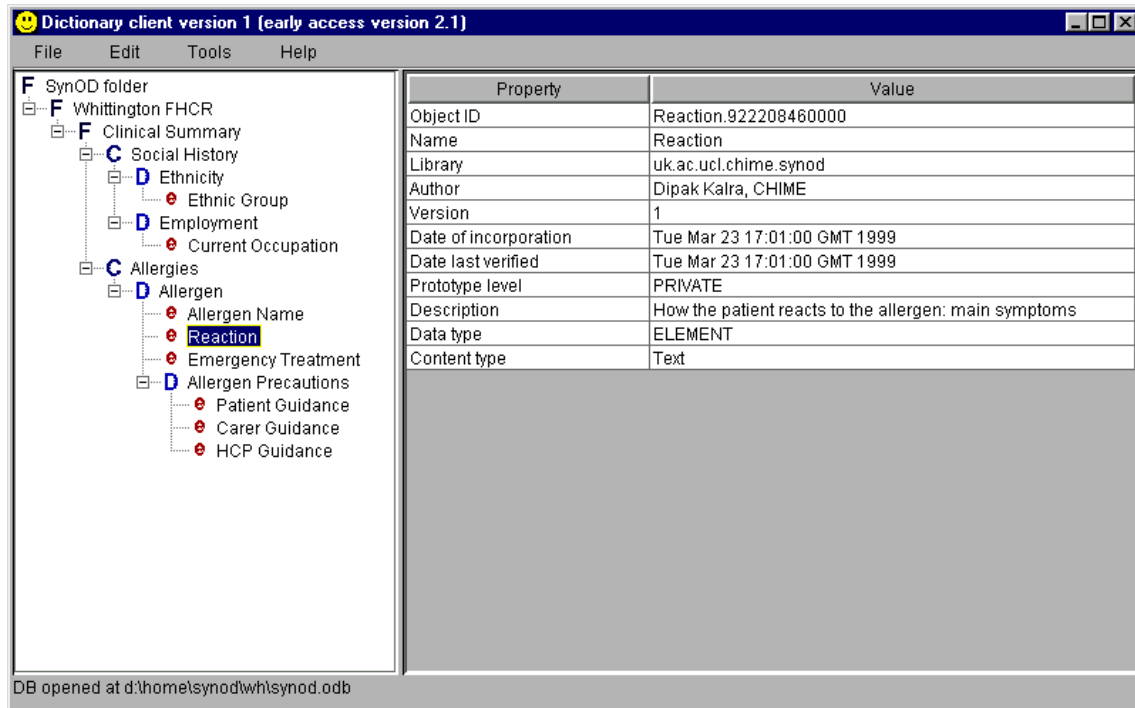
Figure 11: Example Screen from the Object Dictionary Authoring Component

An extended information model for dictionary objects is still under development in collaboration with other consortium members and with other terminology tools developers. The additional features of the dictionary will include:

* synonyms (+ with preferred terms, preferred language, and the configuration of these), synonym-oriented views
* clinical concepts, knowledge tags (with GALEN), CEN standard component name categories (with CEN/TC251 PT27)
* qualifiers, CEN standard component annotations (with CEN/TC251 PT27)
* data entry validation criteria and links to protocols (with ICRF)

The Persons Look-up Service

The UCL Persons Look-up Service is a component providing information on the identification of patients, healthcare professionals and other staff to the other FHCR services. It provides a repository of person names and other demographic information, together with their access rights status, that can be used to identify persons within an EHCR or to authenticate access rights to a given set of record components. The data repository uses and extends Novell NDS objects and its meta-directory, and is accessed via Java Naming and Directory Interface (JNDI) APIs.

Overall System Configuration

Figure 12 below shows the breakdown of tasks in the creation of the anticoagulant system. The 'plug-in' nature of the system is worth emphasising - JNDI provides the ability to plug in data sources, with the medical client utilising a plug-in architecture for the display of anticoagulant control information.

The technical approach adopted at this demonstrator site will inevitably show some measure of specificity to local clinical and technical requirements. However, in documenting the final demonstrator the generic aspects of each component described in this report and of the overall computer engineering environment will be emphasised.
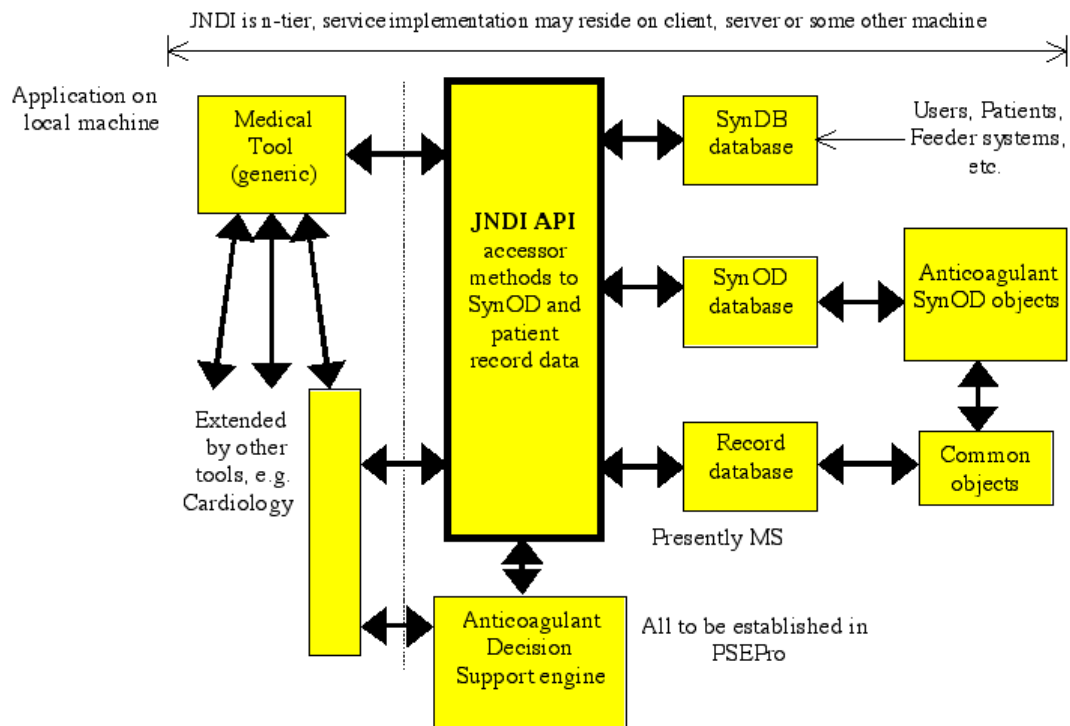

Figure 12: Component Services within the Whittington Anticoagulant System

Post-Project Demonstration: Extended Cardiology Deployment

The core technologies described above place UCL in a good position now to extend the initial scope of the demonstrator. This is vital to provide sufficient profile for the work accomplished to generate confidence in this overall approach to delivering comprehensive and shareable electronic healthcare records.

Funding avenues for the demonstrator are being explored, but some revenue has already been committed by the Whittington Hospital itself, Camden & Islington Health Authority through its Local Implementation Strategy, three industrial partners with an interest in the development of distributed web based record services, and the European Commission though its Fifth Framework IST Programme.

These will together provide the opportunity to continue to build the demonstrator site post-SynEx, and provide an extended marketing opportunity for the SynEx components deployed in it.

The specific strands of further development planned for the medium term are:

- to provide distributed access to anticoagulant records and decision support services to GPs and community pharmacists (expected to commence in October 2000);

- the extension of the cardiology clinical scenario to include angina and myocardial infarction for which both clinical guidelines and established data sets are at an advanced stage;

- collaborative work with General Electric/Marquette to incorporate muti-media investigation reports (in particular, bio-signals) within the federated record;

- component integration work with other consortium members, under SynEx Consortium licences being drafted at present;

- the extension of distributed access to include mobile users, through a new approved Fifth Framework project expected to start in January 2001.

UCL will investigate the potential to utilise the following non-UCL components within the London Demonstrator, through bilateral agreement and/or licences issued at the end of SynEx:

for operational use:
- **ProForma protocol services** (to author new clinical entries for one selected cardiac condition, for care management according to local clinical guidelines)
- **GALEN concept services** (to search the Object Dictionary Client for clinically-related record concepts, to facilitate the authoring of new data schemata and web servlet extraction methods)

for off-line verification:
- **DHE meta-directory services** (to mirror the DHE object schema within the ODC)
- **DHE patient data services** (to enable the DHE to act as a feeder system to the UCL FHCR services)

- **Security components** (the details of the integration of ENATOR and NOVASYS components is being reviewed, awaiting further component details from these participants)
- **TeleNurse and I4C integration**: as candidate clinical applications.

References

1   Bangemann M et al. Recommendations to the European Council: Europe and the global information society. Brussels: European Commission, 1994.

2   Smith R. What Clinical Information do Doctors Need? BMJ 1996; **313**: 1062-8.

3   Audit Commission (UK), For Your Information - A Study of Information Management and Systems in Acute Hospitals. ISBN: 011 886 4165. HMSO Publications, London, 1995.

4   Rector A, Nowlan W, Kay S. *Methods of Informatics in Medicine 1991;* **30**:179-186.

5   Shortliffe EH, Barnett GO. "Medical Informatics: Computer Applications in Health Care". *Addison Wesley Publishing, Reading, Massachusetts.* 1990.

6   The Computer Based Medical Record. Dick R, Steen E, Eds. *Institute of Medicine, Committee on Improving the Patient Record.* National Academy Press, 1991.

7   Deloitte & Touche. Emerging trends in Information Technology. MedInfo Conference Proceedings 8, Part 1, 1995.

8   Kalra D (ed.) The Synapses User Requirements and Functional Specification (Part A). The Synapses Project. Deliverable USER 1.1.1a. EU Telematics Application Programme Office 1996.

9   Southgate L, Berkson L, Fabb W et al. Towards Better Definitions of Competence: a Paper From the Cambridge Conference. Office of the Regius Professor of Medicine. Cambridge University. 1989.

10   The new NHS: Modern, Dependable. Cm 3807. December 1997.

11   Shortliffe EH, Barnett GO. Medical Informatics: Computer Applications in Health Care. Addison Wesley Publishing. Reading, Massachusetts. 1990.

12   Weed L. Medical Records, Medical Education and Patient Care: the Problem Oriented Record as a Basic Tool. *The Press of Case Western University Cleveland, Ohio.* 1971

13   Roper N, Logan W, Tierney AJ. The Elements of Nursing, 2nd Edition. Churchill Livingstone, Edinburgh. 1985.

14   Dick R, Steen E, Eds. The Computer Based Medical Record. *National Academy Press, Washington.* 1991

15   Bishop CW. *MD Computing.* 1991;8:208-215

16   Campbell J, Ginver N, Seelig C, Greer A, Patil K, Wigton R, Tape T. *MD Computing.*1989; 6:282-287.

17   Fischer PJ et al. User Reaction to PROMIS: Issues Related to Acceptability of Medical Innovations. *Proceedings SCAMC.* 1980.

18   Tang PC, Shortliffe EH, Detmer DE. *Annal. Int. Med.* 1991;115:979-981.

19   Lilford R, Kelly M, Baines A. Effect of using protocols on medical care: randomised trials of three methods of taking an antenatal history. *BMJ* 1992;305:1181-1183.

20   Spann SJ. *Journal of Family Practice.* 1990;30:457-464.

21   Shortliffe EH, Barnett GO. Medical Data: Their Acquisition, Storage and Use. *Medical Informatics: Computer Applications in Health Care, Chapter 2.* Addison Wesley Publishing, Reading, Massachusetts. 1990.

22   Pagano MP. Communicating Effectively in Medical Records. Sage Publications, London. 1991.

23   Proposals for New Performance Procedures: a Consultation Paper. General Medical Council, London. 1992.

24   Southgate L, Berkson L, Fabb W et al. Towards Better Definitions of Competence: a Paper From the Cambridge Conference. *Office of the Regius Professor of Medicine.* Cambridge University. 1989.

25   Skifjeld K, Harket G, Hurlen P, Piene J, Skjervold S. A Document Architecture for Health Care Records Integrating Structural and Semantic Aspects. *AIM Office.* 1992.

26   Rodnick J. *Journal of Family Practice.* 1990;30:460-464.

27   Blois M. Information and Medicine: The Nature of Medical Descriptions. *Berkley and Los Angeles: University of California Press.* 1984.

28   Arnold R, Povar G, Howell J. The Humanities, Humanistic Behaviour, and the Humane Physician: a Cautionary Note. *Annal. Int. Med.* 1987;106:313-318.

29   Bryant GD, Norman GR. *New England Journal of Medicine.* 1980;302:411.

30   Hunter K. Doctors' Stories: The Narrative Structure of Medical Knowledge. *Princeton University Press, New Jersey.*

31   Rector A, Nowlan W, Kay S. *Methods of Informatics in Medicine.* 1991;30:179-186.

32   Text of the Council Decision Adopting a Specific Programme of Research and Technological Development in the field of Communication Technologies (1990-1994). Brussels 5/6/91.

33   Decision No. 1110/94/EC of the European Parliament and the Council, adopting a fourth
      framework programme of European Community activities in the field of research,
      technological development and demonstration (1994-1998).
34   Source: http://www.centc251.org/