Healthcare@Home: Research Models for Patient-Centred Healthcare Services

Mahesh Subramanian, Ali Shaikh Ali, Omer Rana, Alex Hardisty and Edward C. Conley
School of Computer Science and Welsh e-Science Centre
Cardiff University, UK
{m.subramanian, ali.shaikhali, o.f.rana, alex.hardisty, e.c.conley}@cs.cardiff.ac.uk

Abstract

Healthcare@Home is a research project implementing a prototype to 'push' or 'pull' data via mobile devices and/or dedicated home-based network servers to one or more data analysis engines. This data is then used to evaluate diabetes risk assessment for a particular individual, and also undertake trends analysis across data from multiple individuals. We outline the need for such a personalized health management system, and describe a logical and physical architecture of the system – making use of the Service-Oriented Architecture (SOA) approach. The current status of the implementation is then reported, utilizing Web Services and clinical sensor technologies.

1. Introduction

Support for individual-centered healthcare is an important trend in modern medicine, primarily aiming to identify how particular treatments and drugs impact individuals rather than a clinical group. This is also a recognition that a "one size fits all" healthcare solution may not adequately address and differentiate between individual-specific needs and particular priorities identified in a national health plan. In this context, the timely identification of individuals who pose a high risk of developing a particular disease has become an important priority. Diabetes is one such disease, the increase of which in adults worldwide is estimated to rise from 135 million in 1995 to 300 million in 2025 [7]. The ability to capture data directly from an individual, and subsequently analyze this for trends that could indicate potential risk to the individual provides a significant advance over current approaches to diabetes management. The increasing availability of mobile devices, the ability to interconnect these devices to each other and the existing wired infrastructure, provides opportunities for seamless patient monitoring [12]. Such monitoring and analysis also needs to be undertaken in a time frame that is meaningful, as the timeliness of prediction is a key challenge for existing systems. In the context of healthcare provision, it is also necessary to account for the reliability of such predictions for computing individualised risk.

1.1. Individualized Healthcare Scenarios

Grid computing involves the coordinated sharing of resources that exist within multiple administrative domains (Virtual Organisations) to solve a single large problem. Recently, Grid computing has been used effectively to support healthcare applications, in areas such as diagnosis, prognosis, disease prediction and drug discovery. Grid computing within these domains has generally been used to support applications which require access to large computational resources or the integration of different types of data (medical images, genomics information, etc) – examples include eDiaMoND [3], MammoGrid [2], VOTES [5], BIRN [4] and BioPattern [1]. The Akogrimo project [6], in particular, focuses on developing an e-Health environment to support medical devices for patient monitoring (e.g. ECG and ultrasonic instruments) and collaboration tools to support interaction between patients, paramedics and medical experts for responding to clinical emergencies. In the Akigrimo scenario, when an ECG abberation is detected, the system first determines the location of the patient, then the location of the responsible emergency dispatch centre and the general practitioner or cardiologist attending the patient. Clinical Virtual Organisations (CVOs) [5] provide a useful abstraction for focusing on the integration of data from a collection of such mobile devices – focusing on security, data access and ownership issues that are pertinent to healthcare provisioning. Grid computing infrastructure provides important services to support such CVOs, especially as queries to such CVOs must span multiple heterogeneous domains and data formats.

The management of Personal Healthcare Records (PHRs) has also been explored by others, and primarily provides an electronic mechanism to integrate healthcare data about a particular individual. A home telecare framework for chronic diseases to collect clinical data from a patient for a "case management team", and enable a clinician to review this data periodically is proposed in [13]. A patient health record, maintained by the system, can be accessed via a low cost workstation – providing health education and a status questionnaire. Medical intervention occurs in response to an observed change in the patient health status by a clinician. The system is constrained by the need for specialist home workstations and a central server to host PHRs. MediCompass [14] is a commercial system similar in concept to the previous framework, but instead allows a user to access their PHR using a standard Web browser. Patients will still need to login into the system and record their medical readings manually. Other commercial PHR systems include iHealthRecord [15] and CapMed [16]. Our system differs from these as follows: (1) patients submit their medical reading automatically via a mobile device. The medical reading is pre-validated prior to storage in a PHR; (2) the system uses data distributed among clinics; (3) the system is integrated with the current healthcare pathway; (4) a trend analysis mechanism is used to compare data across different patients.

1.2. Healthcare@Home Requirements

Healthcare data needs to be timely, of high relevance to an individual's condition and be accessible by authorized members of distributed care teams. A significant challenge in data quality arises when the care pathway continues outside of the clinical environment and its various quality controls. A key requirement to be satisfied by the Healthcare@Home project is to extend this continuum of care to the home environment so that the same range and quality of "timely" data is available to enable better self-management of conditions as well as for clinical management of conditions. Essential requirements that the system has to address in the context of providing healthcare are as follows:

- Ability to collect data from the patient irrespective of the patient's location (clinic, at work, or at home) using biomedical sensors and other clinical equipment. Subsequently, to relay the status of a patient's condition and record this information in a disease-specific "continuing care record";
- Provide opportunities for constructing detailed clinical disease progression models against which the patient's condition can be monitored and treated in an "evidence-based" manner;
- Provision of the right level of information (with appropriate access controls to maintain patient confidentiality and privacy requirements) to both patients and clinical staff, to permit respectively self-monitoring and clinical control;
- Provision of supplementary information (e.g., dietary and other lifestyle information) to the patient to permit more effective self-control of their condition;
- Provide support for risk analysis, permitting appropriate and timely intervention by clinical staff within the patient population e.g. within a specific geographic area.

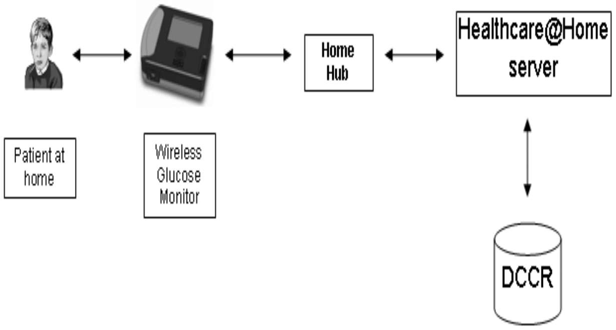

Figure 1. Data Capture Pipeline

In assessing and addressing these requirements we have paid particular attention to the acceptability of the resulting system to its potential end users; clinicians and patients. This is reflected throughout the architectural design and implementation of our research prototype system. Figure 1 illustrates the data capture and analysis pipeline adopted in the Healthcare@Home project. Data is initially captured from an individual using specialist sensors and locally stored within an individuals home on a "Hub" – currently manually generated within the research prototype. This data can then either be pushed to a server periodically, or may be pulled from the server on-demand. Details of the Hub and the structured data are provided in Section 2.2.

2. System Architecture

Based on these requirements we propose a distributed architecture.

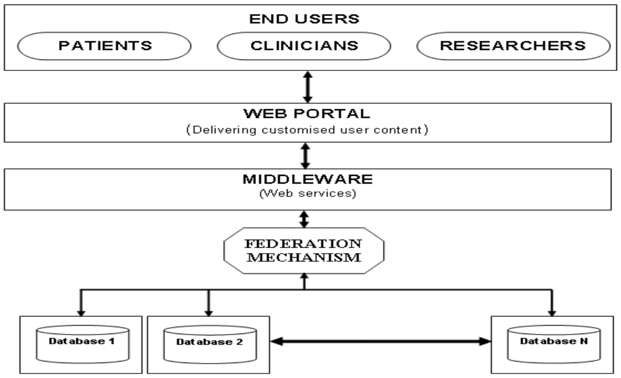

Figure 2. Logical Architecture

2.1. Logical Architecture

The logical architecture is illustrated in figure 2, and provides role-based access to data maintained within multiple databases. Roles within the system include: *Patient:* an individual registered with the system; *Clinician:* which may include a doctor, a member of administrative staff, or a nurse. Each of these clinician roles are treated differently for data access rights; *Researcher*: a user with express ethical permission and authority to access anonymized patient data to study trends, generate statistics etc. Researchers are not given access to data from any individual patient, and can only analyze general trends. Access control mechanisms are associated with each of these roles. It is necessary to have access to patient medical information in all the clinics where the patient may have had treatment. Hence, it is necessary to access distributed data and this can be achieved through database federation.

2.2. Physical Architecture

The logical architecture described in section 2.1 has been realized as illustrated in figure 3. A portal server provides a secure and customizable interface between the pilot end-user (clinicians, patients and researchers) and the middleware. Various portlets have been created for each user role identified earlier, and are all hosted in the same portal server environment. A process server allows data management processes to be deployed and invoked on the server, allowing new processes to be added for additional capability. The process server operates primarily as a workflow enactment engine. The initial request to the process server is generated through an externally accessible Web Service. A database is used for storing patient information in a particular geographic location. A mediator architecture is then used for federating geographically distributed databases. SPINE [9] is the national database managed by the UK National Health Service (NHS) for storing patient information including demographic data. It is part of the NHS Care Records Service – a service focusing on providing electronic access to patient healthcare records in the UK. SPINE is only expected to host summary records, essentially providing a reference point to the location of the full record for each patient. A patient database and schema have been developed, based upon the contents of the NHS Information Standards Board 'Diabetes Continuing Care Reference' recommendation. We refer to this dataset as the DCCR in our system – with the added extension that our DCCR also records sensor data from patients. The PDS (Personal Demographic Service) [10] is a part of SPINE and contains the demographic data for patients. PDS has been replicated locally for this project and includes all the information that would be available in the PDS located in SPINE. Finally, Health Level 7 (HL7) is a standard used in healthcare for the interchange of electronic clinical information among clinical organization. Currently, HL7 v3 is being evaluated by clinics in the UK to interact with SPINE.

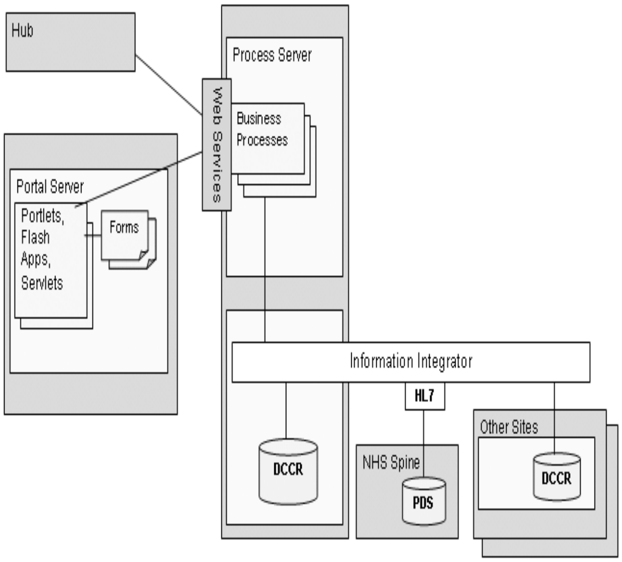

Figure 3. Physical Architecture

As illustrated in the physical architecture, a patient may interact with the system in two different contexts:(i) at a clinic when a patient visits to register, for instance. Initial tests will be carried out in a clinical setting to obtain the initial patient information including demographic and prior medical data. The use of wireless sensor devices is also considered in the clinical setting. Sensor enabled devices are already available for measuring various medical data such as blood pressure, glucose and lipids; (ii) at home, after the procedure in a clinical setting has been carried out, the patients will be provided with a "Hub" and appropriate sensor enabled devices allowing continuous monitoring. This would enable the clinicians to obtain regular medical information (e.g. blood pressure, lipids and glucose) from the patient without a visit to the clinic. The Hub may be used to push patient sensor information to a clinic, and also to provide feedback to the patient. In one scenario, a biometric sensor (finger-print scanner) is used to identify the patient, and a unique"ID" included with the sensor data. The sensors communicate with the Hub using Bluetooth, and the Hub communicates with the Healthcare@Home system through a secure Virtual Private Network (VPN) or HTTPS connection.

2.3. Security Considerations

Security is a significant requirement in a healthcare application involving patient data. Our system uses a two layer authorization scheme: (i) users wishing to access the system must provide their credentials to a portal, which then identifies resources that such a user can access; (ii) each document in the system must have associated credentials and attributes identifying: the actor who created the document and the purpose and type of the document. These credentials and attributes are used to authenticate and authorize the actor who created the document. The use of digital certificates (in place of username-password combination) is being considered for authorizing user access to the portal. Similarly, the finger-print scanner (or a smartcard in the future) is being used in the home setting to identify the patient. In this way, multiple patients living under the same roof, who are registered, can make use of the same set of sensor devices and hub to send their medical data to the "healthcare@home" server. Two separate data stores, the PDS and the patient database (DCCR), are used as a first step to prevent medical data in the patient database being identified to a patient without having appropriate access to both data stores. Full anonymization of data is a formidable research challenge, and currently being considered by a number of other researchers.

The use of Bluetooth to transmit medical information from the sensor devices to the hub in a home setting also has security and power constraints. For secure transfer of data between the hub and the heathcare@home server, an HTTPS connection or VPN tunnel are being considered.

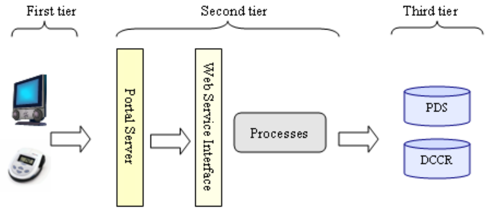

Figure 4. 3-tier Implementation

3. Implementation

The Healthcare@Home system is implemented using a service-oriented approach, consisting of a set of independently implemented services that can be updated or extended on-demand, and consists of three tiers. Our implementation makes use of the IBM WebSphere product family, which includes a portal and process server.The first tier contains access devices or medical sensors that may be mobile or fixed. In most cases, these will be desktop computers or tablet PCs running standard Web browsers. Each user would have access through a portal with access constraints determining the set of portlets that they can make use of. The second tier provides a Web Service interface, and a predefined workflow for managing data obtained from the user. The third tier consists of medical and demographics data repositories used to keep a persistent record of the data. Figure 4 illustrates the different tiers in the system. The managed information consists of:

- Personal details: includes but is not limited to demographics such as name, date of birth, and address. Personal details are provided by the national patient record service.
- Care relationships: details of those providers who are involved in an individual's care. Examples include doctors and nurses.
- Health events: which arise from care "encounters" an individual has with the healthcare system, and may reflect minor experiences through to life-changing occurrences. Health events of importance are reflected in current clinical communications, for example initial visits, referrals, and some requests and reports of tests and procedures.
- Sensor events: which are generated when a particular reading is made on an individual. This reading may be automatically captured by an implanted sensor, by the patient using a personal measurement device, or by a clinician in a hospital/laboratory.

- Health status: current view of care provided to an individual, for example the current medication that has been prescribed for them.

Figure 5 illustrates the physical implementation of the system. Various components of the system are illustrated, with the DB2 (UDB v8.2) database used to implement the DCCR and PDS, an IBM x336 server used to host the Websphere Portal Server, an IBM x236 server to host the Websphere Process Server, and a Blade centre for managing eight JS20 blades. Each JS20 blade runs the DB2 database to host a local DCCR – this is used to simulate multiple clinical domains. QUantitative Individualized Risk Analysis (QUIRA) is used to analyze data from a particular individual, and compare this with data across a population, to assess risk factors that are most likely to impact on disease progression for a given individual. The QUantitative Individualized Diagnostics modelling environment (QUID) is used to support data integration from a variety of sources – such as from sensor devices (via the hub), from data entered into forms, and data obtained from other specialist instruments (such as retinopathy cameras).

3.1. System Interactions

In our system, health or sensor events detected in tier 1 lead to the generation of a "document" – this term is used in accordance with the definition provided by the IHE [11], and refers to information that can be grouped based on context (such as time it was produced) or content (such as information relating to a particular patient, a clinician, etc). These documents may be generated as a result of six possible external interactions with the system by a user (a patient, clinician or researcher – defined in section 2.1):

- Patient *pushes* data to a server: occurs when the patient sends his medical reading to the server via a hub.
- Patient *pulls* own data from server: occurs when a patient requests his medical information to be displayed.
- Clinicians *pushes* patient data into system: occurs when a clinician registers a patient or sets clinical targets for a patient.
- Clinicians *pulls* patient data from server: occurs when a clinician requests a patient health record.
- Clinicians requests patient data to be *pushed* to patient: occurs when a clinician sets intervention points for an individual patient.

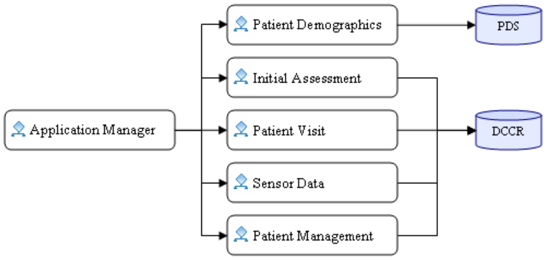

Figure 6. Process Interaction

- Researcher *pulls* aggregate data from server: occurs when a researcher undertakes trend analysis on the recorded data of multiple patients.

To make the system scalable, we support two types of push mechanisms: (i) a "continuous" mechanism that allows all monitored data to be submitted from tier 1 to the server or from tier 2 to the clinician or patient; (ii) a "triggered" mechanism that allows a subset of the total data to be submitted to the server, based on a policy decision made by a clinician and a system administrator.

3.2. Process Management

The process architecture is illustrated in figure 6, with the "Application Manager" acting as a process coordinator. The remaining processes are responsible for dealing with a particular health event. For example, the "Initial Assessment" process is responsible for handling the patient initial assessment visit. The processes are implemented in BPEL and the data type for the inputs and output is text.

3.3. User Interaction

A JSR168 compliant portal server, with portlets implemented using Java Server Pages is used, enabling access to the healthcare data via different user roles. JSR168 defines a portlet specification, including a contract between the portlet container and the portlet. The data formats of wireless (and other) devices can be built into the middle tier to control the amount of information transmitted to the client devices; subsequently this information is marked-up for display on that device. Five portal pages have been implemented, with each page reflecting a phase in the diabetes healthcare pathway – each page can have multiple portlets. These pages are: "Registration", "History", "Measurement", "Management" and "Dashboard". The registration page handles patient registration in the demographic service and enables patient registration at the

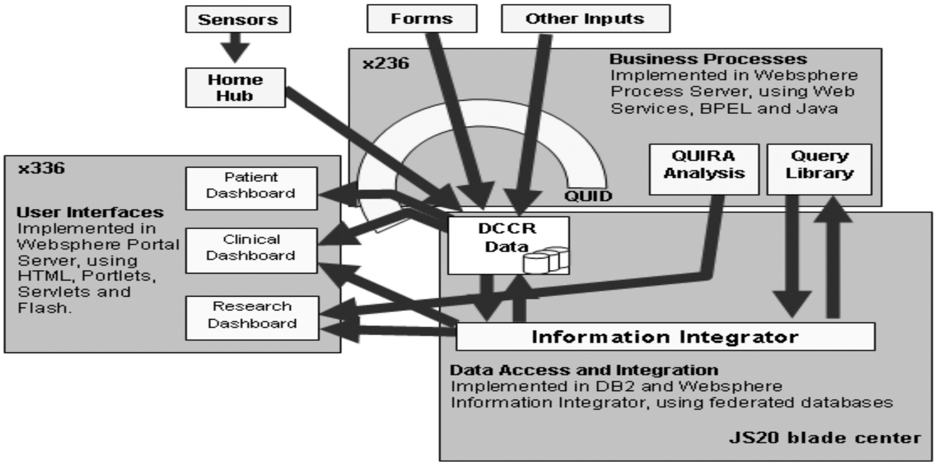

Figure 5. System Implementation

local diabetes clinic. The history page involves recording of lifestyle information (such as whether the patient smokes). The measurement page allows data recorded at a clinic to be entered into the DCCR. The management page allows a clinician to set targets on blood pressure, glucose and lipid content, to enable these targets to subsequently be compared with the measured values. The dashboard enables clinicians to visualize the patient's data. Each portlet can be considered as a separate application. The portlets within different pages can exchange data with each other. This is achieved by initializing a portal session for each user – subsequently sharing session parameters between all the portlets. The session stores the parameters that need to be passed between the portlets.

The portal provides a single sign-on capability, via username/password pair. For instance, in the clinical role only the patient selection portlet is enabled. Once a patient has been selected, the patient identity is stored in the portal session, and other portlets are enabled. Adobe Flash forms technology is used to create a dynamic user entry form, and for data binding between the form and the Web Services. The portlet passes the patient identity to the form, resulting in a SOAP call to our Web Service to retrieve relevant patient data. The returned data (in XML format) is displayed in the Flash form.

4. Conclusion

A research prototype system to support personalized health trend analysis is presented, with monitoring of parameters that are primarily aimed at individuals with diabetes. Grid computing technologies are used as a basis to integrate data about a particular individual from a variety of different sources. For instance, an individual may potentially visit a variety of specialist clinics over their lifetime. Grid computing infrastructure is used to support integration of different types of data (such as sensor data directly from a patient and image data from retinopathy), along with support for trends analysis between data from different individuals. The trend analysis aspect is a key requirement for a "proactive" public health management programme, where near real time status information can be used to prioritize disease management, and to identify recommended actions needed to reduced individualized risk. Similarly, the use of portal technologies plays an important role in making the outcomes from such data analysis available from a number of perspectives to various types of users.

Acknowledgement: We would like to thank Steve Wright and Andrew Borley of IBM Hursley Laboratory, UK – also part of the Healthcare@Home project – for their contribution to the architecture and implementation. Other partners include Zarlink Semiconductors, Smart Holograms and the Diabetes Research Unit at Cardiff University.

References

[1] L. Sun, P. Hu, C. Goh, B. Hamadicharef, et al., "Bioprofiling over Grid for eHealthcare", Proceedings of HealthGrid conference, Valencia, Spain, 2006.

[2] D. Rogulin, F. Estrella, R. McClatchey, "A Service-Based Grid Environment for Managing Distributed Medical Images", 6^{th} International Conference on Web-Age Information Management, 2005.

[3] S. Lloyd, M. Jirotka, A. Simpson, R. Highnam, D. Gavaghan, D. Watson and J. Brady, "Digital Mammogra-

phy: a world without film?", Methods of Information in Medicine, Vol:44, No.2, 2005.

[4] V. Astakhov, A. Gupta, J. Grethe, E. Ross, D. Little, A. Yilmaz, M. Martone, X. Qian, S. Santini, M. Ellisman, "Semantically Based Data Integration Environment for Biomedical Research", Proceedings of the 19^{th} IEEE International Symposium on Computer-Based Medical Systems, 2006.

[5] R. Sinnott, A. Stell and O. Ajayi, "Development of Grid Frameworks for Clinical Trials and Epidemiological Studies", Proceedings of HealthGrid conference, Valencia, Spain, 2006.

[6] C. Loos, "E-Health with Mobile Grids: The Akogrimo Heart Monitoring and Emergency Scenario", EU Akogrimo project Whitepaper, 2006. Available for download at: `http://www.mobilegrids.org/`.

[7] H. King, R. Aubert, W. Herman, "Global burden of diabetes, 1995-2025: prevalence, numerical estimates, and projections" Journal: Diabetes Care, Volume: 21 Issue: 9, pp 1414-31, 1998.

[8] Health Level 7 standard. Available at: `http:/www.hl7.org/`. Last accessed: June 2006.

[9] SPINE Information can be obtained from: `http://www.connectingforhealth.nhs.uk/publications`. Last accessed: June 2006.

[10] Personal Demographic Service (PDS). Details at: `http://www.connectingforhealth.nhs.uk/publications`. Last accessed: June 2006.

[11] "Integrating the Healthcare Enterprise", IT Infrastructure Technical Framework, Vol:1 Augest 15, 2005. Available at: `http://www.himss.org/ASP/topics_ihe.asp`.

[12] V. Shnayder, B. Chen, K. Lorincz, T. Jones, and M. Welsh, "Sensor Networks for Medical Care", Technical Report TR-08-05, Division of Engineering and Applied Sciences, Harvard University, 2005

[13] B. Celler, N. Lovell, J. Basilakis, F. Magrabi, M. Mathie, "Home Telecare for Chronic Disease Management", Proceedings of the 23^{rd} Annual EMBS International Conference, Istanbul, Turkey, 2000.

[14] MediCompass PHR. Available at: `http://www.imetrikus.com/about.asp`. Last accessed: June 2006.

[15] iHealthRecord. Available at: `www.ihealthrecord.org`. Last accessed: June 2006.

[16] CapMed. Available at: `www.capmed.com`. Last accessed: June 2006.