Towards a Personalised Health System

Peter Weller[1], Alberto Fernández Gil[2] and Eduardo Alsonso[3]

[1]*Centre for Health Informatics, City University London, London, U.K.*
[2]*CETINIA, University Rey Juan Carlos, Madrid, Spain*
[3]*Department of Computer Science, City University London, London, U.K.*

Keywords: Clinical Decision Support Systems, Personalised Healthcare, Software Agents.

Abstract: This paper presents a Personalized Healthcare System (PHS), a decision support tool that can adapt to changing conditions, such as aging and illness, in individual patients. The system consists of three components: a unique personalised profile, a collection of web based tools and a web based repository for managing interactions between clinicians and tools. The proposed system makes extensive use of software agents, both for collecting the initial information required to construct a personalized profile and for transporting the information needed to use the on-line decision support tools. The paper discusses the operation of a PHS and suggests possible implementation issues.

1   INTRODUCTION

This paper introduces the concept of a Personalised Healthcare System (PHS). It discusses basic components, describes a number of functions of a PHS and offers some suggestions on implementation.

A PHS is a software-based system that is unique for every member of a population. Depending on the patient, on their current health status and location, it could either be running continuously (as a semi-active or active Decision Support Sys-tem (DSS)) or activated only when required (as a passive DSS). A PHS changes and evolves over the life of the associated person by adapting to events such as aging, illness, accidents and life-style choices. Furthermore, it can link with the PHS of family members, friends and neighbours to gain knowledge of hereditary or location-based diseases or illnesses.

The purpose of the PHS is to support the care and treatment of individual patients in a variety of situations, primary and secondary care, and telehealth via tele-consulting. The system includes access to both diagnosis tools for assessing current health status and prediction tools to consider "what-if" scenarios for the effect of treatments and life-style changes.

Healthcare has been one of the target application areas of expert decision support systems for many years, including recent distributed technologies such as multi-agent systems (Paranjape and Sadanand, 2009 & Shirabad et al., 2012), and SOA (Nadkarni and Miller, 2007). However, to the best of our knowledge, frameworks supporting the integration of different DSS and generic profile knowledge for health care has not been addressed by previous works.

2   STRUCTURE OF PHS

The PHS is composed of three main components:
1) A client program, which constructs a Personal Patient Profile (P3) and co-ordinates the information flow process for each patient;
2) A central web-based repository that stores basic pro-files and manages interactions between clients and remote tools.
3) A collection of remote web-based tools for DSS and modelling/prediction tasks;

2.1   Client Program

A core task of the client program is to construct the P3, a key component of the PHS. The P3 is constructed prior to any consultation from three resources, namely, the patient's Electronic Patient Record (EPR) EPR, a generic patient profile, and the latest information from external data sources such as news links, journals and social media sites. The

client program also contains the mechanisms for sending requests via soft-ware agents and receiving information back from other soft-ware agents.

2.2 Central Repository

The central repository contains four elements:
1) A library of generic profiles for a range of patient scenarios;
2) A database of decision support tools;
3) A database of modelling tools;
4) A suite of management tools for organising and up-dating the previous elements.

Generic Profiles

The library of generic profiles consists of sets of data relevant to different patient scenarios, ages and conditions, for example, a four-year-old girl with diabetes or a fifty-year-old man with high blood pressure.

The base profile used will change during the patient's life. When a baby is born it is given a profile for a new-born baby. A range of these will be available for different initial conditions such as premature births, multiple births or birth complications. Some facts in the initial profile could be obtained from the parents' and siblings' EPR. As the baby ages, the profile changes to include expected conditions such as measles and chickenpox, and developmental features, for example, teething, walking and the beginnings of speech. If a condition or feature is no longer relevant to the baby it is not included, so milk teeth will not appear in a profile for a fifteen-year-old male. Over time, the requirements for making decisions related to baby health are replaced by those associated with young children. These, in turn, are superseded by those for older children and then for adolescents, young adults, etc. In addition, details for illnesses, accidents and events not generally considered part of the basic health scenario, for example the development of diabetes, are added as required. Eventually, the profile will contain information associated with aging and geriatric conditions, such as dementia, rheumatism or senility. Finally when a person dies the PHS is retained as a source of information for descendents, friends and neighbours. New users joining the system could be given an initial profile based on sex, age and basic health.

Database of Decision Support Tools

This database contains the details of all decision-support tools registered on the system. The information stored includes: a brief description of the application domain; input data required for operation; cost of service (if applicable); a measure of accuracy and expected results from the developer (in terms of sensitivity/specificity), and finally feedback for users both as a rating and as text.

Database of Modelling Tools

Following the same approach as the one outlined above, this database stores details of modelling and prediction tools stored in the system.

Suite of Management Tools

The management tools perform two tasks: sending and receiving information via intelligent agents, and the maintenance of the database of third party decision support and modelling tools.

2.3 Web based Tools

Decision Support Systems

This component of the system consists of a suite of individual DSSs that are developed and maintained by independent third parties (e.g., universities, charities or disease specific research groups). They are registered with the PHS central repository and made available via the internet. Each module could be developed for a specific function, for example the diagnosis of an illness or condition. These modules would be constructed using a variety of data analysis, decision-making or prediction tools selected by the developers.

Modelling Tools

Similarly, this component consists of a suite of modelling tools (MTs) for predicting a range of scenarios. They are also developed and maintained by third parties and registered on the PHS central repository.

3 OPERATION OF A PHS

The following scenario illustrates a typical exchange between a GP and a patient, although the methodology would be very similar for a clinician caring for a patient in a hospital. Before each consultation, once the patient confirms their attendance, the client program constructs the personalised patient profile (P^3). To do this *Agent A* is sent to the Central Repository with brief patient details. *Agent B* is returned with the most appropriate generic profile. This profile is then integrated with the patient's EPR and any relevant news items or pertinent issues from social media resources. Finally, intelligent agents are used to conduct an on-line search for any relevant new

treatments and evidence-based guidelines. These steps are shown in Figure 1.
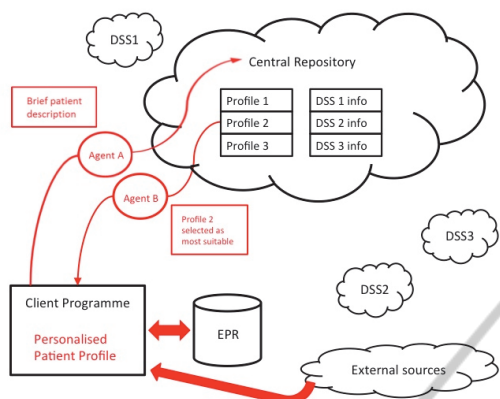

Figure 1: Selection of Profile and construction of P³.

Figure 2 depicts the initial situation, with a patient meeting with their GP. The GP starts to assess the patient's condition with a set of questions. During the consultation the patient's symptoms and responses are added to the P³ (and into the EPR) via speech recognition or typing. These inputs are used to refine any search strategy to provide more relevant and more recent guidance.
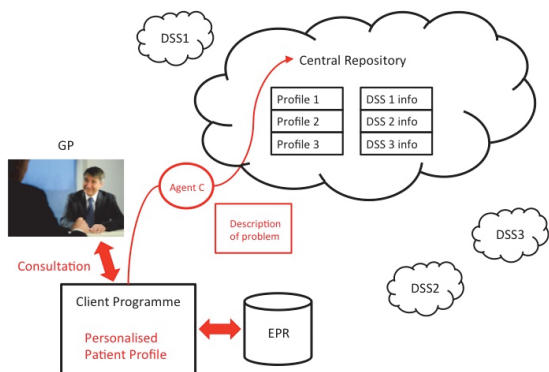

Figure 2: Meeting with GP and request for DSS support.

From the information and data collected during an initial dialogue a preliminary diagnosis would be made. In addition, central resources may need to be consulted for advice or clarification. In order to do so, the program despatches *Agent C*, loaded with some basic individual information extracted from the EPR and from the consultation (for example, fifty-year-old male, chest pains, hypotension) to the central repository. This stage is shown in Figure 2.

The central repository now matches the information supplied to a list of available DSSs using the descriptors provided by the DSS developers. Not all DSSs will be relevant to the current problem. Only those that match the problem details (to a greater or lesser extent) are returned to the co-ordinator with a list of the inputs and information required for the DSS, such as a feedback rating. This situation is illustrated in Figure 3, where two DSSs (DSS 1 and DSS 2) are identified as possibly being suitable for the requested task. This does not mean that each identified DSS is necessarily able to perform the required diagnosis, just that there is some match between the problem and the capabilities of the DSS. For example, DSS 1 may have been developed for diagnosing tumours in a fifty-year-old male, whereas DSS 2 may have been developed for diagnosing a heart attack. The information collected by the coordinator is sent to the clinician via *Agent D*.
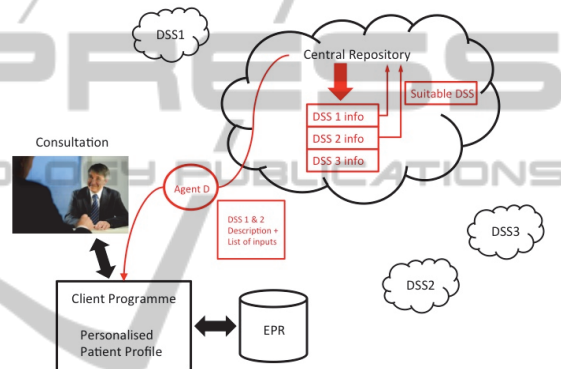

Figure 3: Selection of possible DSS and dispatch of details about suitable DSS.

The clinician now decides which DSS is most capable of diagnosing the patient's condition. This could be based on a number of criteria, including relevance of problem domain, data required as inputs for the DSS and even a ranking based on previous performance. There could be a cost function associated with using a DSS, for example a privately developed DSS could be made available but with a charge levied each time it is selected. In Figure 4, DSS 2 has been selected and *Agent E* is sent, with the required data, to the central repository in order for the diagnosis to be reached.

The central coordinator then sends *Agent F* with the data to DSS 2. This approach maintains confidentiality and ensures that no patient identifiable information is sent to the DSS. This stage is shown in Figure 5.

Once the diagnosis has been reached an agent is sent back to the central coordinator with the output. The data is also added to the DSS knowledge-base for further development of the tools. In Figure 5, *Agent G* has been dispatched by the Central
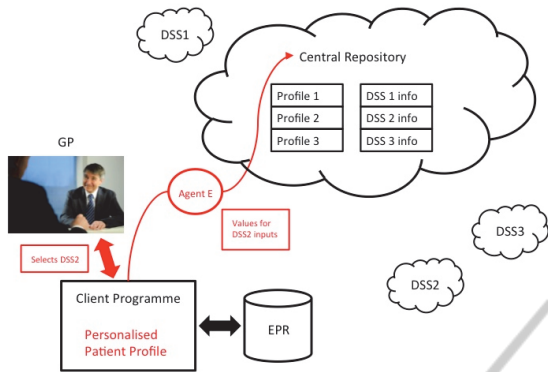
Repository with the output from DSS 2.

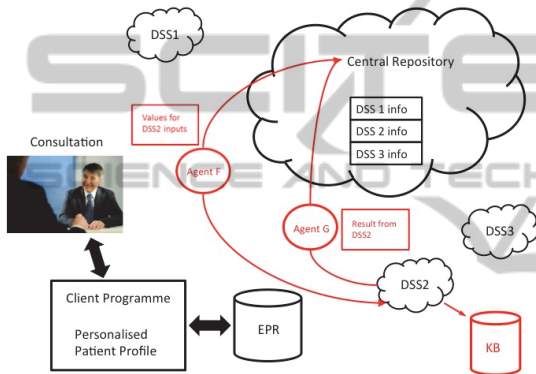Figure 4: Selection of DSS and data for diagnostic request.

Figure 5: Transmission of data to selected DSS and transmission of DSS result.

The central coordinator then transmits the result to the GP through *Agent H*, as illustrated in Figure 6. This information can be added to the patient's EPR for future reference. Finally the Client Program sends *Agent I* with the GP's feedback on the performance of DSS2. This is added to the central register for use with future requests.

The figures above only show the main tasks. Additional support-functions such as receipt of information, safe transmission and encryption are not included but would be components of the process nevertheless. A similar approach would be adopted for modelling the effects of a treatment or drug regime.

4 SYSTEM IMPLEMENTATION

In this section we present some ideas on the implementation of the PHS. We consider service-oriented architectures (SOA) and ontologies as the enabling technologies towards the implementation of
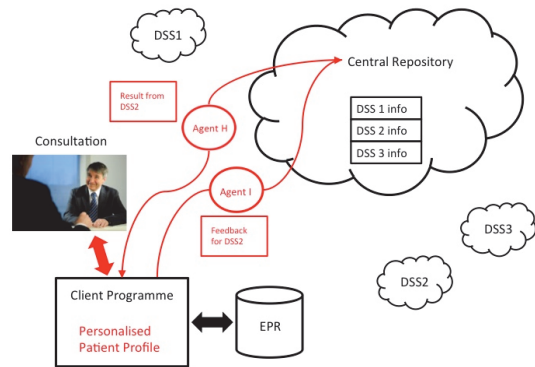

Figure 6: Transmission of DSS output and feedback.

this kind of systems.

SOAs (Huhns and Singh, 2005) comprise at least two components: *service providers* and *service clients*. Services are software components that encapsulate some functionality. Providers and clients interact for some service to be carried out. Basically, the client provides the inputs to the service (e.g., patient symptoms), and the provider returns the results/outputs of the execution of the service (e.g., diagnostic). Typically, a third component, a *service registry* (also known as a *directory*) is also present. Providers advertise their services by registering a description with the directory. An example of information included in a service description can be the type of inputs expected and of outputs provided. Service descriptions can be specified at different levels of expressivity ranging from purely syntactic to complex logic-based descriptions (Fernandez et al., 2012). An *ontology* is a specification of a conceptualization (Gruber, 1993). Ontologies are used to share information/knowledge, or more specifically to share the vocabulary used when agents/services interact. For example, it is important that the concept "blood pressure" is shared, and then understood the same way (e.g. units), both by the GP and the diagnosis tool.

Ontologies are used in PHS for knowledge representation such as generic profiles, EPR, P^3, DSS/modelling tool descriptions and communication message contents.

Figure 7 depicts the main building blocks of the PHS architecture. As presented in previous sections, the PHS is composed of three main elements, a client program, a set of remote tools and a central repository. We propose the use of a service-oriented architecture (SOA) to coordinate the interaction among the different actors in our PHS.
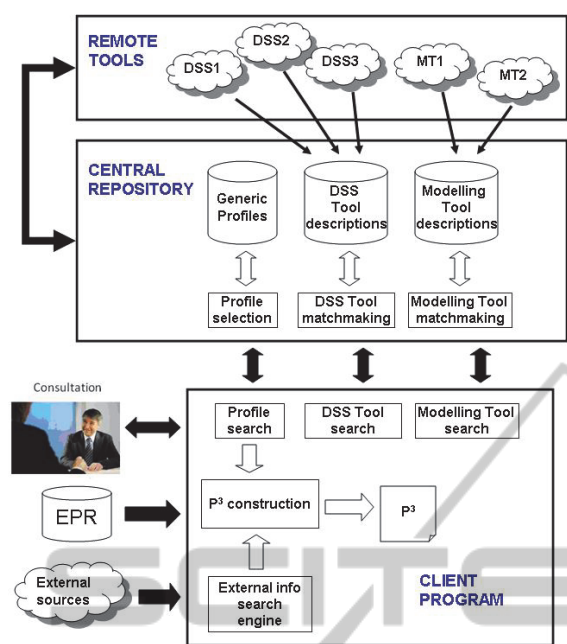
Figure 7: PHS architecture.

In the rest of this section some details about each component are given.

4.1 Remote Tools

DSS and modelling tools in the PHS are represented as (web) services in our SOA. Providers advertise their services by registering a description of their tools with the *Central Repository*. The set of existing DSS and MTs is not fixed and can vary over time. In fact, it is natural that new tools appear as time goes. Tools are created and maintained by third parties. Different tools may differ on their specialty, generality/specificity, technique, cost, etc. Several tools might target the same objective (e.g. disease) but be provided by different institutions/companies. In some cases, aspects such as cost, trust and reputation might have to be considered so as to select the appropriate provider. The proposed SOA architecture gives flexibility to developers to add, modify or remove tools. They only have to register the description with the *Central Repository*.

4.2 The Central Repository

The main functionality of the *Central Repository* is to store and manage basic patient profiles and remote tools. Thus, it contains three databases, namely generic profiles, DSS tool descriptions and modelling tool descriptions.

Each *generic profile* comprises two elements, a

profile pattern description (e.g. four-year-old girl with diabetes) and the profile detailed information. Profile descriptions are used to identify the adequate profile according to the characteristics of the patient provided by the GP through the *Client Program*. This task is carried out by the *Profile selection* module. Concept similarity (Euzenat and Shvaiko, 2007 & Fernandez et al., 2007) techniques can be applied for this task.

The *Central Repository* also provides matchmaking functionalities to locate DSS and modelling tools (*DSS and Modelling Tool matchmaking* modules in Figure 13). Tool providers register a (semantic) description of their DSS tool (e.g., diagnosis of lung cancer in teenagers) with the *Central Repository*. When the GP (or just a user) decides to consult a DSS they specify a description of the desired service. The *Central Repository* then matches the request against the registered tools and returns the information about the ones appropriate for that task. Descriptions for these tools are more complex than the ones for patient profiles so advanced methods will be used here. In particular, techniques from the semantic web service matchmaking (Klusch, 2008) will be adapted and applied for the specific particularities of DSS and modelling tool discovery.

Note that, although conceptually centralized, a distributed implementation could be adopted for efficiency or fault tolerance. For example, three different repositories each containing a database and its corresponding matchmaking functionality would be a straightforward distribution, but others would be also possible.

4.3 The Client Program

The *Client Program (CP)* is in charge of mediating between the GP and the rest of the system, which includes the construction of the P^3 as well as the interaction with the *Central Repository* to find generic profiles, DSS or modelling tools. The construction of the P^3 is the main and more complex task carried out by the CP. As previously described, it requires the EPR, a generic patient profile, and information from external data sources. Ontologies have been proposed as a technique for automatic processing and interoperating healthcare representation standards like SNOMED CT[1], HL7[2], OpenEHR[3] or CEN 13606[4] (Gomez-Perez et al.,

---

[1] www.snomed.org
[2] http://www.hl7.org/
[3] http://www.openehr.org/
[4] http://www.cen.eu/

2009; Sahay et al., 2012 & Schloeffel et al., 2006).

The *External info search engine* module is in charge of accessing external information sources (databases, journals, social media, etc.). Ad-hoc wrappers might be needed for different sources. We believe that the increasing adoption of Linked Data[5] as a standard way of exposing information will ease the integration of the different information sources. In fact, there are already several related linked data sources such as DailyMed[6] or MediCare[7], which provide information about existing drugs.

5 CONCLUSIONS

In this paper we have proposed a Personalised Healthcare System (PHS). Key characteristics of such a system are its uniqueness for every member of a population; its capacity to evolve over the life of the person; and its flexibility to coordinate external resources such as DSS/modelling tools and information sources. We described the operation of a PHS and gave some ideas towards its implementation following a SOA approach.

It is part of our future work the implementation of a PHS as well as extending the architecture improving aspects such as mechanisms for tool selection (e.g. trust and reputation, auctions, etc.).

ACKNOWLEDGEMENTS

REFERENCES

Euzenat, J., Shvaiko, P., 2007. *Ontology matching*. Springer, Heidelberg (DE).

Fernandez, A., Cong, Z., Balta, A., 2012. Bridging the Gap Between Service Description Models in Service Matchmaking. *Multiagent and Grid Systems*, 8 (1): 83-103.

Fernandez, A., Polleres, A., Ossowski, S., 2007. Towards Fine-grained Service Matchmaking by Using Concept Similarity. *Workshop on Service Matchmaking and Resource Retrieval in the Semantic Web (SMR2 at ISWC)*, pp. 31-45.

Gomez-Perez, J. M., Kohler, S., Melero, R., Serrano-

Balazote, P., Lezcano, L., Sicilia, M.A., Iglesias, A., Castro, E., Rubio, M., Buenaga, M., 2009. Towards Interoperability in e-Health Systems. *International Conference on Health Informatics*: 205-210.

Gruber, T. R., 1993. A translation approach to portable ontologies. *Knowledge Acquisition*, 5(2): 199-220.

Huhns, M. N., Singh, M. P., 2005. Service-Oriented Computing: Key Concepts and Principles. *IEEE Internet Computing*, 9(1):75–81.

Klusch, M., 2008. Semantic Service Coordination. In *CASCOM - Intelligent Service Coordination in the Semantic Web*. Birkhaeuser Verlag, Springer: 59-104.

Nadkarni, P. M., Miller, R. A., 2007. Service-oriented Architecture in Medical Software: Promises and Perils. *Journal of the American Medical Informatics Association*, 14 (2), pp 244-246.

Paranjape, R., Sadanand, A., 2009. *Multi-Agent Systems for Healthcare Simulation and Modeling: Applications for System Improvement*. ISI Global.

Sahay, R., Zimmermann, A., Fox, R., Polleres, A., Hauswirth, M., 2013. A Formal Investigation of Semantic Interoperability of HCLS Systems. *Interoperability in Healthcare Information Systems: Standards, Management, and Technology*. IGI Global.

Schloeffel, P., Beale, T., Hayworth, G., Heard, S., Leslie, H., 2006. The relationship between CEN 13606, HL7, and openEHR. In *Health Informatics Conference*, Sidney, Australia.

Shirabad, J. S., Wilk, S., Michalowski, W., Farion, K., 2012. Implementing an Integrative Multi-agent Clinical Decision Support System with Open Source Software. *Journal of Medical Systems*, 36 (1), pp 123-137.

---

[5] http://linkeddata.org/

[6] http://www4.wiwiss.fu-berlin.de/dailymed/

[7] http://www4.wiwiss.fu-berlin.de/medicare/