

Information-centric Pervasive Healthcare Platforms

Dirk Trossen
Non-affiliated
Colchester, U.K.
dirk_trossen2000@yahoo.com

Dana Pavel, Kenneth Guild
School of Computer Science
and Electronic Engineering
University of Essex, Colchester, UK
dmpave, kguild@essex.ac.uk

Jean Bacon, Jatinder Singh
Computer Laboratory
University of Cambridge
Cambridge, UK
jmb25, jatinder.singh@cl.cam.ac.uk

Abstract—Pervasive healthcare is gaining attention outside the research community. But real-life deployment of many solutions faces the realities that come with everyday life, namely the existence of boundaries between applications, networks and devices. These often impair the end user experience through failing along these boundaries or enforcing particular end user experiences at them, e.g., by providing single-device or single-network solutions. We intend to address these shortcomings with an approach that places information at the centre of the platform design, including the low-level communication framework. We believe that this approach alone can address the various governance and provenance conflicts that emerge in real-life deployments. This paper describes early work by outlining our design vision as well as early findings on its realization.

Keywords – Pervasive health, information networking

I. INTRODUCTION

Future healthcare is currently an important issue on various agendas, from local to global organizations and governments, from industry to academia. The number of people living with long-term health conditions is growing as people's lifespan increase, placing a heavy burden on healthcare systems and support networks. Data from the World Health Organization shows 75% of the total population with one chronic condition and 50% with two or more conditions [1]. The risk factors for chronic diseases are often related to lifestyle choices like smoking, alcohol intake, physical inactivity, or poor diet. The lifestyle choices have such an influence that some of these chronic diseases are now called "lifestyle diseases" and prevention is getting more important with every day. One way of realizing prevention is to empower people by creating self-monitoring systems that can give them enough knowledge to understand the effect that their actions have on their well-being and support them in making better lifestyle choices. Prevention can make a real difference for the low risk category of patients, around 70-80% in the UK [1], meaning that with the right support they can learn how to manage their disease and avoid or delay further deterioration.

Monitoring and communications technology have developed to the extent that they are potentially exploitable in large-scale, widespread healthcare. At the same time, people are becoming increasingly familiar with and accepting of a wide range of technologies. The time is right to consider technology-assisted lifestyle monitoring that could help us detect and prevent various chronic diseases.

Although technology is continuously improving, we are in danger of fragmentation in healthcare provisioning if the current approach of vertically separated applications proliferates. We need solutions that provide a framework that allows for proper data recording, analysis, retention, usage and auditing across the boundaries at application, network and device level. Furthermore, we need solutions that take into account users' contexts and policies, that use such information for adapting not only applications and services but also communication and networking layers across domains such as networking, organizational and even device domains. Hence, we need solutions that operate *across* domains, rather than forcing end users into them, to make systems work, e.g., through providing solutions that only work on a single device or within a managed single network.

Our contribution to this area is the development of a multi-domain system that incorporates patients at home and outside, primary care practices, hospitals and outsourced services such as X-rays. To realize the operation within a realistic multi-domain environment, our work is guided by an information-centric vision, i.e., information, its structures as well as its governance, is preserved throughout all layers of the system. This enables communication layers to be aware of information collected through applications and vice versa. The work presented here is currently developed as part of the PAL project¹. This paper will mainly focus on our vision centered on a lead scenario that reflects challenges as well as solutions we intend to address. For this, Section II starts with the lead scenario and the main challenges introduced by it. Section III describes a high-level system design that focuses on main components and their roles in our vision. In Section IV, we present our findings based on our current experiences in developing such a system followed by conclusions in Section V.

II. SCENARIO AND CHALLENGES

The following is intended as an example scenario to highlight potential issues and challenges that will be faced when deploying such solutions in real life:

Tim is 55 and had heart surgery 2 months ago. The operation was a success and he is recovering well. The doctor advises him to take it slowly and pay more attention to his lifestyle. To make it easier, Tim is using the PAL system, which can record various activity data from personal

¹ The work for this paper has been funded by the UK TSB/EPSC project PAL (<http://www.palproject.org.uk>).

devices and correlate it into an easy to understand presentation so that he can see how certain activities impact his physiological state, especially his heart condition. Tim used to have quite a hectic schedule; he would skip meals and not drink enough, adversely affecting his blood pressure. PAL helps him better understand such relationships, and some of his data also helps his doctor find possible causes when something does not go as expected. One day, during shopping, he almost fainted. The PAL system, which also includes an on-the-go panic service, was able to immediately alert the medical team, which in turn could consult his personal data to see what he was doing before the incident, as well as gain access to other medical data available through his GP and the hospital where he had his operation.

There are several key challenges to be addressed when realizing this scenario. While they are not particularly novel in the context of healthcare scenarios, we focus on them throughout the paper, highlighting that *information* is central to implementing the functionality, and we use this focus as a cornerstone for our design.

The first challenge relates to *gathering* information. Information comes from diverse sources and includes data from various sensors, activity information derived from emails, computer logs etc., to wider community knowledge. This poses the challenge of obtaining information in a consistent way that enables application modules to properly operate on it. Second, the gathered information needs to be *provisioned*, capturing aspects of creation of service relations, temporary and permanent storage and definition and enforcement of appropriate policies. Third, is *processing* the information but also adapting the operation of the overall system to perform efficiently and correctly according to the desired functionality. Fourth, raw but also processed information needs to be *presented* to various end users, e.g., Tim and the medical team, on various devices in a meaningful way. Finally, all of these operations of gathering, provisioning, processing and presentation need to work across several application, network and device domains in ways that maintain the end user experience as seamlessly as possible. This touches upon the issue of *governance* of operations and their underlying information.

III. AN INFORMATION-CENTRIC SYSTEM DESIGN FOR A PERVASIVE HEALTHCARE PLATFORM

We address the above challenges of real-life deployment through a design that places information at its heart. This includes *all* layers of the platform, in particular the networking and communication functions. Although it is difficult to present the breadth of this design, and its impact on dedicated components, in the space available, the organization of the paper, following the challenges presented above, is intended to aid this presentation.

A. Main system components

Figure 1 presents a high-level view of our system. While the layering of components: user; application (the lifestyle assistant); middleware and communication levels is a traditional design, the integration of information, its structures and its governing policies across *all* levels is the

differentiating criterion. That is, information structure, not its particular semantics, is preserved from the user level down to the communication and routing levels.

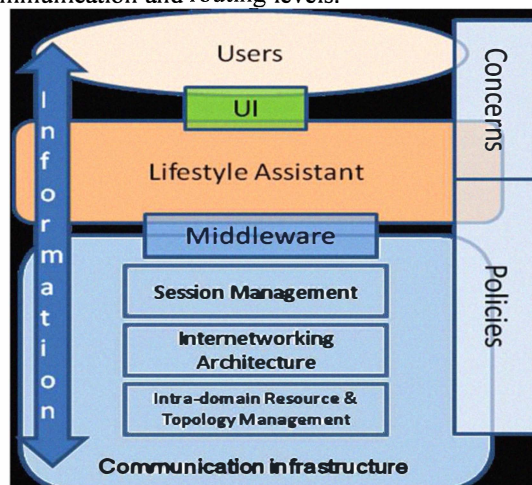


Figure 1. High-level component view of the system

The following sections describe the system components following the challenges of gathering, provisioning, processing, presenting and governing information.

B. Information gathering

Our design integrates various data sources being collected within our scenarios. These consist of multiple types of data from various user devices, such as mobile phones, laptops, PCs, and body sensors, producing information like a user's context (e.g., location, movement, ECG, heart rate, social context, sound, etc.), messages, calendar events or media. We also integrate existing knowledge bases into the system, such as the *NHS Direct* system [5]. To enable seamless services, the gathering of network capabilities is important, such as available networks, link failures etc. While this is not used directly by end users, it is important for the network awareness of the system.

This results in a gateway-based design, in which sensing *gateways* encapsulate the specifics of the particular data to be integrated, such as the implementation of a wireless sensing protocol or access to a particular database. The gateway is responsible for mediating the semantic differences with respect to the PAL system, resulting in the first step of information processing, namely the alignment of information semantics. Data provided by the gateways is then provisioned and processed in the following steps.

C. Information provisioning

Once data has been gathered via gateways from various sources, information is provisioned within the PAL system for processing and presentation. Since healthcare is a highly data-driven environment, we provision information as *events* that are relevant to other components and encapsulate *health incidents*, including observations, such as sensor readings, doctors' notes or a perceived emergency; as well as actions, such as pressing the panic button.

Architecturally, this provisioning is largely implemented in the *middleware*, a software layer that resides between

applications and the network. Acting as a layer of abstraction, the middleware provides communication mechanisms that operate independently of specific application/network concerns. As such, it provides a suitable point for policy enforcement. While such middleware interacts directly with the application, it resides on a communication infrastructure for transferring information between devices by virtue of a variety of networking technologies. Our vision is to build the middleware directly on an information-centric routing infrastructure that works differently from the current Internet, i.e. characterized by a *publish/subscribe* rather than a send/receive service model. Furthermore, information is directly addressed, not the communicating endpoints via a flat label identifier. This preserves higher-level information structures within the communication infrastructure without relying on the information semantics being known. With this, we provide an asynchronous, many-to-many, information-centric routing paradigm that is particularly suitable for the event dissemination middleware described above. With this, we achieve a scalable decoupling of clients, i.e., as clients communicate through the middleware, they are not burdened with the addressing specifics of every potential information source/sink. However, the information labeling in the communication layer allows for an efficient mapping of event labels onto routable information addresses.

Finally, the chosen pub/sub service model of the middleware and communication layers allows easy integration into a database/storage system. Typically, health information requires persistence; so a coupled infrastructure removes the overhead of separate messaging and distribution substrates. Further, the powerful data management functionality of database systems provides a rich representation of state from which event transmission (routing) and disclosure policy can leverage.

D. Information processing

Information is not only provisioned in its original form but is further processed to aid presentation to the user. For this, we use aggregation techniques for well-known input sources, such as location and activity. But we also integrate the end user directly into this process by allowing annotations to gathered data, thus harvesting the user's knowledge. In addition, available knowledge base sources, such as NHS Direct, are integrated into the processing step, i.e., gathered information is aggregated and interpreted through wider community knowledge. This leads to a growing pool of information that incorporates real-life sensing, user knowledge about certain events, and common knowledge from an expert community.

E. Information presentation

Collected data enables various visualizations to be created, helping users understand causal relations between their actions and their wellbeing. Our intention however is not to build diagnostic tools but to support users in reflecting on their activity history as well as, assuming users' consent, to help healthcare providers find missing information.

The scenario described above has a strong emphasis on self-monitoring and self-understanding. All the collected data, even when incomplete, enables us to create a more complex picture of what the user experienced during the day. The presentation of information is centered on *meaningful events* for the end user. In a dialogue with the user, the PAL system determines the specific meaningful events that the user cares about. It then creates a visualization that walks the user through this series of events along a set of selected *stories*. The preservation of information structures throughout the system, down to the communication infrastructure, aids the construction of appropriate stories and the integration of appropriate events in the visualization.

F. Information governance

Healthcare information is inherently sensitive. Therefore, information must be protected at all levels of the system. Clearly, the patient must be able to describe the circumstances in which particular information is disclosed to particular parties. Such policy must be context-sensitive, as restrictions will often vary according to circumstance. A motivating example is an emergency situation, where restrictions may be relaxed (see Section II).

Patients will undoubtedly have *specific* concerns, e.g., requesting location being generally hidden, or that a GP cannot access data unless explicitly authorized. However, it is unrealistic, if not dangerous, to expect a patient/user to define all situations in which disclosure is appropriate. This is because the typical patient will neither have the knowledge nor expertise to precisely determine the information relevant to each party. In practice, we expect that a disclosure policy will be influenced by providers, possibly through the use of policy templates describing particular interactions. Any involved party is implicitly responsible for protecting health information. In line with this, any specific patient preferences must be adhered to by everybody involved.

Given that assisted living scenarios operate outside of a single domain of administrative control, it becomes necessary to consider issues of distribution. In addition to controlling the disclosure of information to particular clients, it is necessary to consider the intermediate nodes through which information is passed. Further, multiple clients will often be interested in the same information. As such, it becomes necessary to consider the controlled 'broadcast' of information, ensuring the appropriate level of disclosure as relevant for each recipient in the group. The chosen combination of an information-centric middleware operating on top of an information-centric routing infrastructure provides the leverage of enforcing disclosure policies deep in the infrastructure, across various domains and involving several parties. This extends to the issue of accountability, which is important given the sensitivity of personal health information. In addition to persisting the health data itself, information detailing the context (circumstances) surrounding a disclosure must also be recorded. This enables appropriate investigation in situations of concern, allows those responsible for information to demonstrate that they have met their obligations, and can assist in the discovery of any issues in the set of governance policies. A coupled

storage-messaging middleware facilitates audit in the sense that access and transmission of data can be recorded with the data itself. This accountability extends further to network operations, such as session continuity and network selection. The preservation of information structures on the communication layer allows for creating accountability of these network functions in the light of, e.g., malfunctioning. For instance, in our scenario Tim could account for the costs involved in his PAL system through associating network costs with particular parts of his monitoring activity.

IV. CURRENT FINDINGS

Work is underway in implementing and preparing the early deployment of design solutions. This is based on existing and currently expanded infrastructure that was devised within the HIPNet project [3]. The project established a heterogeneous networking environment that connects the main project sites at Essex University, Cambridge University and BT Research, using optical as well as all-Ethernet networking technologies. There is, in addition, wireless coverage through WLAN and WiMax technologies available at Essex University that can be leveraged to produce a truly multi-network environment to verify our design. Mobile as well as fixed device technology is part of the environment, being used for sensing as well as presentation of information. For instance, the mobile phone based NORS platform [4] is used for gathering phone-based as well as worn sensor based information.

Our design work, as well as the early realization in an actual networked environment, has led to a number of findings. First, we came to realize that the integration of a large number of information sources does not add much to the overall system evaluation. Instead, we tried to create a multi-device and multi-network mixture that is representative for deployment problems in many scenarios. With this in mind, we focused on mobile devices as well as available activity information on fixed computing devices for the information sources being utilized in our scenario.

Second, since we base our design on information visibility throughout *all* layers, it becomes apparent that the current internetworking solution in the Internet, namely the Internet Protocol (IP), is not suitable for our approach. IP provides a bit transfer service to (information-aware) applications with information being opaque to the underlying layers. Hence, the objectives of optimizing resource usage, e.g., at the optical transport layer, only come with expensive cross-layer operations due to the opaque nature of IP. As a result of this, the PAL project integrates results from the PSIRP project [2], including a novel information-centric internetworking protocol. This offers a simple and scalable model for information identifiers over which the underlying routing and forwarding infrastructure operates. Initial work is underway to integrate this work into our environment to fully utilize its potential in our design.

Third, the aspect of information governance throughout the system is crucial for our selection of scenarios. While many scenarios for healthcare are placed in the traditional formal healthcare area, we chose to focus on scenarios in the informal healthcare area, i.e., targeting lifestyle management

and similar scenarios. This forces us to involve the end user and her preferences in the reconciliation of potentially conflicting policies. Assuming a higher freedom of end users to choose their policies, the reconciliation problem becomes harder. Furthermore, the multi-domain character of user scenarios, highlighted in Section II, increases the multi-party character of this reconciliation process, in particular, parties from various societal domains (compared to the rather healthcare-focused nature of traditional healthcare scenarios). In our scenario of Section II, Tim's private governance policy in addition to the medial team's formal policy needs reconciliation during the execution of the scenario. This can easily be expanded when integrating family and friends into the informal care circle.

While the presented findings are largely qualitative, they turned out to be important ones to formulate the research agenda of our design and its implementation. We intend to complement these findings with dedicated quantitative findings on the level of, e.g., user perception, efficiency of resource usage, policy compliance and the like.

V. CONCLUSIONS

Pervasive healthcare systems are facing many real-life deployment challenges that are hardly ever addressed in the design of these systems. This often leads to disruptions in end user experiences or constraints being imposed when using these systems, such as single device or single network solutions. Overcoming these challenges is crucial in making headway in further adoption of future healthcare systems.

This paper outlined an information-centric design for a pervasive healthcare system. We believe that such an approach will enable us to overcome the challenges caused by conflicting governance policies for resources and information being used throughout a scenario. This is achieved by exposing the information, its structures and its governance policies throughout the entire system, including the lowest level communication platform. This enables gathering, processing, provisioning and presenting information in a holistic way.

Evaluating such a design and its associated claims of better performance is a challenging task in itself. For this, the PAL project has been established not only to develop system solutions but also to model and put in place an environment that enables evaluation. Early findings have led us to make choices on scenarios and components that make us confident that we will provide this evidence.

REFERENCES

- [1] UK Department of Health, "Improving Chronic Disease Management," Report, 2004.
- [2] EU FP7 PSIRP project available at <http://www.psirp.org>
- [3] UK TSB/EPSRC HIPNet project available at <http://gov.epsrc.ac.uk/ViewGrant.aspx?GrantRef=EP/E002382/1>
- [4] D. Trossen, D. Pavel, "NORS: An Open Source Platform to Facilitate Participatory Sensing with Mobile Phones," International Conference on Mobile and Ubiquitous Systems: Networking and Services (MOBIQUITOUS 2007), August 2007.
- [5] NHS Direct available at <http://www.nhsdirect.nhs.uk>