Personal and Social Communication Services for Health and Lifestyle Monitoring

Jean Bacon[1], Jatinder Singh[1], Dirk Trossen[1], Dana Pavel[2], Andreas Bontozoglou[2], Nikolaos Vastardis[2], Kun Yang[2], Sarah Pennington[3], Sarah Clarke[3], Glyn Jones[3]

[1]Computer Lab
University of Cambridge, Cambridge, UK
{jean.bacon, jatinder.singh, dirk.trossen}@cl.cam.ac.uk

[2]School for Computer Science and Electronic Engineering,
University of Essex, Colchester, UK
{dmpave, abodoz, nvasta, kunyang}@essex.ac.uk

[3]Thales UK Research and Technology
Reading, UK
{sarah.pennington, sarah.clarke, glyn.jones}@uk.thalesgroup.com

Abstract—**The focus of the PAL project is to study how future requirements for healthcare services impact current and future communication infrastructures. Current assisted living scenarios work only in restricted environments and are unable to provide continuous user support. Systems in this space tend to be closed, providing only particular functionality and/or operating on specific infrastructure. Our aim is to deploy appropriate interface and interaction paradigms, regardless of the underlying networking and software technologies, to enable users to achieve self-monitoring and self-management of their lifestyle and health. We also consider the integration of social and professional support networks, while ensuring appropriate access controls. This paper details a layered architecture for supporting a diverse range of care services. We provide an integration framework to demonstrate the practicality of our approach, and present results highlighting design considerations for user-centric care systems.**

Keywords-self-monitoring; lifestyle; healthcare; privacy; security; communications; networks; middleware

I. INTRODUCTION

The number of people living with long-term health conditions is growing as people's lifespans 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. Lifestyle choices have such an influence on health that prevention of disease through self-monitoring is becoming increasingly important.

Self-monitoring systems give people additional 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]; with suitable support and engagement they can learn how to manage their conditions and prevent 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 to help in detecting and preventing chronic diseases.

Despite technological improvements, we are in danger of fragmentation in healthcare provisioning if the current approach of developing vertically separated applications continues. In other words, solutions tend to operate in "vertical silos". As a consequence, they are typically restricted to a single domain, for example, working on a single device or within a managed network. Instead, solutions are required that operate across domains, be they organisational, networking, or even device specific. Open systems solutions must adapt to a variety of applications and services, through communication and networking layers that operate across domains. Within our project, entitled *PAL (Personal Assisted Living)*, we envisage a multi-domain system comprising patients at home and outside, primary care practices, hospitals, and outsourced services (such as X-rays and emergency response teams), where components are used and reused to meet the functional requirements across the range of care-related applications and services.

To make such a vision manageable we have started from a number of use cases that highlight the requirements motivated by the above: (1) Prevention of disease onset or deterioration through lifestyle and health monitoring; (2) Integration of social and professional networks; (3) Mechanisms for information governance; (4) Dynamic coordination of system components to realise functional goals; (5) Operation across organisational and communication domains. Section II presents a lead scenario to exemplify these issues, outlining the main challenges for developing a holistic systems solution for healthcare provisioning.

We address these challenges through providing an integration framework, so that a range of applications can be deployed above a middleware layer under which a variety of communications technologies can operate, while supporting diverse device technologies. From the perspective of information management, our framework supports data recording, analysis, retention, usage and audit across the boundaries at the application, network and device level. From the (patient) users' perspective, their data is sensitive and private and we support the specification and enforcement of security policies. Moreover, we envisage the involvement of social and professional networks of carers, particularly as users become less able to cope alone. Controlling data access becomes more subtle as more users (roles) become involved.

All of these aspects are realised in the PAL architecture presented in Section IV. It is this architecture and the presentation of its on-going realisation that is the core contribution of this paper. We recognise that the envisioned system and its early prototyping in a lifestyle management setting is only a first step towards an integrating framework for the healthcare application industry. We believe, however, that our work can provide useful insights in the relevant areas of development. In addition to our integration approach, we also outline experiences with certain aspects of a working system that we built based on our approach. We outline important aspects regarding the system components that we implemented and present first results from user experiments.

We begin with a scenario to highlight the challenges in provisioning healthcare systems, and follow with related work. Section IV details our architecture, and in Section V we describe the realisation of our framework and present results illustrating design considerations for user-centric systems.

II. Lead Scenario and Challenges

Let us first present our lead scenario in the area of lifestyle management. We will use this scenario to highlight the various challenges that exist for realising a system-wide solution for healthcare provisioning.

A. *Lead scenario*

Oscar is a patient-user who has chronic heart disease and needs to pay attention to his wellbeing. His related blood pressure issues make him susceptible to fainting. Since he was diagnosed, Oscar started using the PAL system to collect information about his daily activities, and physiological data relevant to his condition. PAL provides Oscar with support within a *preventative* mode as well as a *reactive* mode. In the preventative mode, the system allows Oscar to track his wellbeing by interpreting and visualising collected data in personalised ways. It also allows Oscar to share certain information with his support network (e.g., family, friends or healthcare specialists). In the reactive mode, the system can provide emergency support, e.g., during an fainting episode.

B. *Main system challenges*

Self-monitoring is all about information pertaining to an individual. We therefore divide the system challenges into the dimensions of *gathering and providing* information throughout the system, *processing and using* the information, as well as *governing* the information according to defined policies.

When gathering and providing information throughout any self-monitoring system, it is important to recognise that assisted living data may be collected from a wide variety of sources. Multiple types of data may originate from various user devices, such as mobile phones, laptops, PCs, body sensors and environmental sensors. The information used may include location, movement, ECG, heart rate, environmental context (e.g., noise levels, temperature, weather, pollen levels, etc.), social context, messages, calendar events, and many more. In addition, available knowledgebase sources, such as NHS Direct [3], could be integrated into the processing step, i.e., gathered information is aggregated and interpreted through wider community knowledge. Each data source may play a role in meeting a number of functional goals, across a range of applications; for instance, an ECG sensor may provide information relevant to applications for a patient-user when exercising, their GP for analysis/diagnostic purposes, and for paramedics in dealing with an emergency. The relevance of information to applications will vary, depending on the circumstances. As described in our lead scenario, we consider two usage modes: *preventative* and *reactive*. Due to the different requirements, the preventative mode focuses more on storage, high-level processing, visualisation, and policies for sharing and remote access. Given the plethora of information being gathered in the preventative usage mode, understanding the information is particularly challenging for system users. Hence, any visualisation needs to go beyond traditional graph-based approaches to achieve the awareness that our users need.

The reactive mode is mainly concerned with reconfiguring data streams at runtime, respecting the various control policies as well as the heterogeneity of the underlying infrastructure. For the latter aspect, it is particularly important to avoid restricting any solution to single domains but rather to work over a variety of available communication technologies as they happen to be available. This requires that the supporting communication infrastructure facilitates information exchange between a wide range of data sources and sinks in a dynamic, context-aware, and multi-domain manner. For that, knowledge of network capabilities, such as availability, link failures, etc., is crucial. While this information is not used directly by end users, it is important for the provisioning of seamless services, for example when users move from indoors to outside or when patients move between areas with variable connectivity.

Clearly, heterogeneity places a heavy burden on any application that is directly exposed to these issues. Hence, abstracting the communication specifics through a common interface is important, aiding the re-use of functional components across use cases. The level of abstraction that a healthcare application shall be built upon is that of information flows that can be controlled in terms of gathering, processing and usage, all under a common policy framework that defines the security properties of the information flow. Such a policy governance framework is important, given the personal nature of the information considered in our scenarios. While encryption of information is important, governance extends to context-aware policies defining access to flows of information (even if encrypted), forming appropriate communication relations (e.g., between doctors and care personnel), and the ability to override standard policy (e.g., in cases of emergency).

We now present a brief overview of related work before outlining our architecture, focusing on information gathering, provisioning, processing, presentation and governance.

III. Related Work

Self-monitoring solutions of different kinds have been well studied in academia and corporate research during the past years, with areas such as health and telemetry [9][10] being well represented within the commercial space. Some examples are movement and fall detection systems, such as WristCare

[11], SenseWear BMS [12], Philips LifeLine [13] and the Wellcore Emergency Response System [14], or heart monitoring solutions such as t+ blood pressure [15], HealthBuddy [16], and CardioNet [17]. For most of the existing solutions, data is stored on a remote/central server and analysed by professionals; in some cases patients might be able to add certain notes, symptoms, and so on, and receive feedback from medical staff. Mostly, the systems are integrated into a purpose-built device, while mobile phones are merely used as an interface device, to input data or to receive alerts.

With the availability of sensors such as the Alive Heart and Activity Monitor [18] that can be connected via Bluetooth to recording devices, such as mobiles or laptops, mobile devices can also be used for providing self-monitoring solutions. For example, MobiHealth provides a suite of physiological data recording through mobile phones [19], where various sensors can connect to the phone and use it to transmit data to a server. As smartphones become increasingly popular, and already include sensors such as GPS and accelerometer, fitness and assistance applications are appearing (e.g., iFall for Android [20], Sports Tracker for Symbian [21], or Endomondo [22]).

Common to all these related approaches is the purpose-built nature of their solutions. What is lacking is a system approach that addresses our identified challenges from Section II through a coherent and holistic design. The following Section IV outlines our contribution to providing such a system solution.

IV. PAL SYSTEM

A. System Architecture

The PAL core architecture follows a traditional layered structure where the middleware masks details of the underlying communications by presenting a communication interface to components that run above it. Although such a middleware can support applications directly, we have developed an additional lifestyle assistant layer capable of providing a higher-level interface to support users during self-monitoring in preventative mode. Other applications can coexist with the lifestyle manager by working directly above the middleware layer, so the infrastructure is generally suitable for supporting pervasive healthcare environments. For example, during reactive mode, external components are coordinated to respond to an emergency situation, as later described.

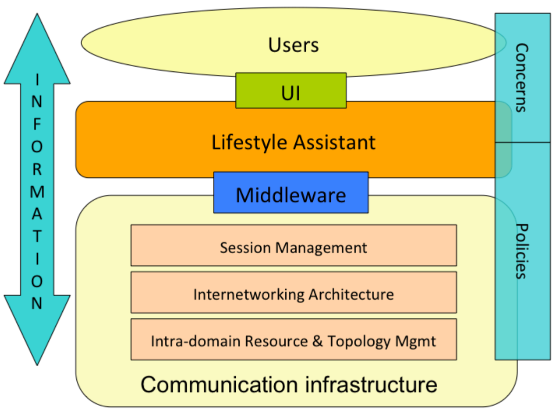

Figure 1. High-level component view of the system.

Figure 1 presents a high-level view of our system. In addition to the traditional layered components—the application (the lifestyle assistant), middleware and communication layers—we highlight the integration of information, its structures and its governing policies across all levels, from the user level right down to the communication and routing levels. Based on this integration framework, we provide in Section V an example of a working system that we realised alongside our lead scenario of lifestyle management. In addition to the realisation of our framework, we also outline first experiences that end users had with this system.

B. Information gathering and provisioning

In our system, much data is provided by sensors—physical (e.g., body or environmental sensors), or virtual (e.g., activity information derived from calendars, social interactions, application usage)—within a heterogeneous environment (see [2][7][7] for more information on gathering capabilities).

To assist in managing the diversity of sensor-produced data, PAL uses 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, and/or access to a particular database or converting low-level sensor data into a format that can be understood by other components in the system. A gateway is therefore responsible for a first level of filtering and abstraction of information. In practice, the scope of each gateway will depend on the particular sensors involved: some gateways might operate to manage, aggregate and transmit the data of a number of different sensors and/or sensor networks, while others may be comparatively lightweight, e.g., tied to a particular sensor, such that it appears as though the sensor directly interacts with other components in the system. Once data has been gathered from various sources, information is provisioned for processing and presentation. Data must flow to a range of different applications and components in order to achieve particular functional goals. Since healthcare is a highly data-driven environment, we provision information as events that are relevant to other components. An event can encapsulate health-related information, including observations such as sensor readings, doctors' notes or a perceived emergency; as well as actions, such as the press of a panic button. Clearly, events will be relevant to a range of different applications/components, as appropriate to the particular circumstances.

Middleware provides a layer of indirection between applications/services and the underlying infrastructure. The role of the middleware is to facilitate interoperability, to provide a reliable mechanism for the exchange of information between system components and services. In PAL, we extend *StreamBUS* (*SBUS*) [4], which is a middleware for managing the exchange of messages (encapsulating events/data) between system components. SBUS was specifically designed for managing streams, enabling dynamic system reconfiguration (connections/disconnections and altering privileges), and providing content-based *publish/subscribe* semantics. However, it aims to be all-encompassing, by also supporting more traditional *client-server* *(RPC)*-based interaction paradigms (see Figure 2). This is important, as communications middleware supporting pervasive healthcare environments must provide for a number of interaction types, to enable a wide range of application-level functionality.

Each application/service in the system becomes an SBUS component, leaving the middleware to manage the communication concerns on its behalf. Data is encapsulated in typed messages (representing events), which are transmitted to other components in the system. The middleware ensures type-safe communication, and enables content-based filters to be imposed on communication channels. From an architectural perspective, much of the provisioning is implemented in the middleware using a *publish/subscribe* semantic, meaning that communication encapsulates the events themselves.

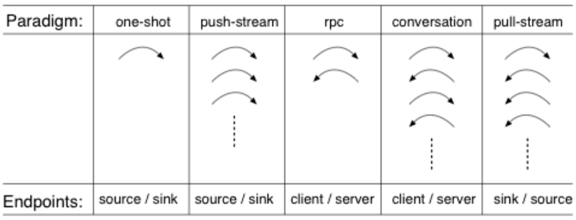

Figure 2. SBUS Interaction Paradigms.

Middleware is a natural point for policy management and enforcement, as it is independent of specific application and network concerns. From the middleware perspective, policy concerns dynamic reconfiguration in response to particular changes in circumstance: specifically, initiating/terminating data flows and/or changing privileges. In environments, such as assisted living, that consist of many data sources, forcing each application to initiate and manage its interactions quickly becomes infeasible (hence the application-specific silos of today). It follows that management increasingly concerns *coordination*: where particular circumstances (events) determine the interactions between applications. Policy operates to dictate how/when particular system components should (or should not) interact to meet the functional goals of a range of applications. In our infrastructure, we have developed a policy engine to manage these concerns. We implement a policy engine within database infrastructure (PostgreSQL), where policy is encapsulated in a set of rules (triggers). A rule is (automatically) executed on a particular change in context, which can operate to: a) raise an alert; b) effect a middleware reconfiguration (e.g. alter a privilege or connection); and/or c) change the active rule set (e.g. to enable a different set of governance policies to apply in an emergency situation).

As illustrated in Figure 1, immediately below the middleware layer is the PAL *communication infrastructure*. It is composed of two key functional sub-systems: session management and inter-networking. Session management ensures that an assisted living application session (e.g., an ECG stream) is kept running continuously when the application device switches from one wireless network to another (e.g., from an indoor WiFi to an outdoor 3G base station) due to user mobility and changes in network availability. The ultimate goal is to provide always-on connectivity to mobile clients, if/when required. For this, we consider an environment where different types of wireless networks co-exist. Typically, when the point of attachment changes, the device's address will also change, and as a result the application session is dropped. To solve this problem, a medium-independent gateway (MIG) has been designed to switch between different types of networks while ensuring session continuity. The essence of the MIG is a tunneling technology which hides any network interface (Layers 1, 2) and IP (Layer 3) changes from the upper layers, thus providing a virtual always-on available gateway. Such a tunnel enables the mobile device to physically join networks that belong to different administration domains without interrupting any on-going connections/flows. This method effectively simulates the behaviour of Mobile IP (MIP) [5] where a Home Agent (HA) is used instead of a tunnel end-point. An algorithm is implemented that carries out network selection based on an individual network's status such as signal strength, bandwidth, energy consumption, cost, etc.

While the MIG is more concerned about heterogeneous *wireless access* networks, the PAL communication infrastructure also addresses *wired core* networks. In particular, a clean-slate approach to internetworking called Information Centric Networking (ICN) [6] is adopted as an IP alternative. Its key task is to conduct route calculation and packet forwarding that directly operates on information items rather than endpoint addresses. The ICN software used is called Blackadder [23], which was developed in the context of the EU project PURSUIT [6]. Blackadder has been slightly tailored in terms of network interfaces and integration with the MIG to accommodate the PAL requirements. Specifically, we developed a link-aware *topology manager* (TM) that makes use of network link status information to assist *route calculation*. The calculated routing paths connect the sources of the information (publishers) with the destinations (subscribers) utilizing the underling publish-subscribe semantic of ICN. The key contribution here lies in enabling *differentiated dissemination* of information; a requirement that we see as important in future healthcare environments. For instance, in cases where certain information, such as heart rate, is deemed more critical than others, an information-centric approach to infrastructure allows for differentiating this specific information compared from other, less important, without resorting to cost-intensive solutions such as deep packet inspection (DPI). In addition, the focus on information dissemination in contrast to endpoint connectivity (as in the current IP-based Internet) allows for realising scenarios where applications are primarily concerned with receiving particular information rather than connecting to a *specific* source for this information. For instance, an application can simply request heart rate data rather than discovering the particular sources (and the specifics that are concerned with contacting these sources) that might potentially provide heart rate data. Hence, we believe that such an information-centric infrastructure approach greatly enhances the capabilities of application developers through providing an information-centric interface for information-rich use cases.

C. *Information processing and presentation*

In pervasive care environments, information will be relevant to a number of different applications and system components, that use data in a variety of ways, e.g. for processing, aggregation, storage, presentation, etc. Sometimes information can be provisioned in its original form (as produced by the information source), but often it requires further processing in order to fit its intended purpose.

In our preventative scenario, further processing of the information collected is vital so that Oscar can more easily understand what happens during each day based on data collected. A different kind of processing is required when certain information is shared with other parties within his support network. For the PAL lifestyle assistant application, we use context type-based information processing, which groups the various sensors into several main categories: *emotional*, *social*, *availability*, *mental* (*interest*), *activity*, *physical*/*spatial* and *temporal*. Within each context dimension, we use a rule-based system that can convert recorded data into high-level information based on various expert or individual experience-based knowledge bases. Additionally, the patient-users can provide their own interpretations either during the recording (i.e., through phone widgets) or later, at the time of visualisation.

Collected data forms the basis for various visualisations created to help users understand causal relations between their actions and their wellbeing. Such visualisations present information at various levels: detailed (e.g., timeline graphs), summarised (e.g., word clouds for text-based data), and correlated abstractions. The last type is a novel type of visualisation we have developed inspired by human storytelling [7]. For that, we process information along a story line by bringing all available (and processed) information together in order to create a summarised (and more engaging) view of the important events during a day. Figure 3 shows an example of an event within a daily story.

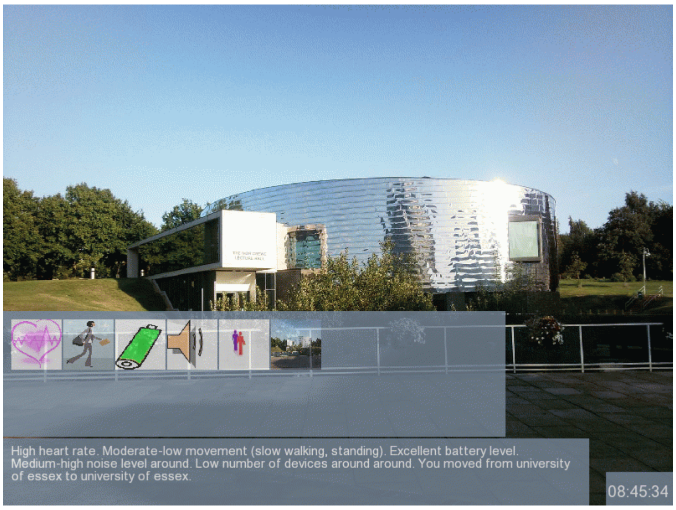

Figure 3. Example of an event within a story.

By providing multiple levels of detail within visualisations, we created a system that can support various types of user interactions and purposes. For example, while a story can show the main events during a day, the user (either as patient or healthcare specialist) can use detailed visualisations to focus on specific aspects, such as heart rate variations over time, social interactions over time and in certain locations, and so on.

Within the reactive scenarios, the coordination of applications and services underpins information processing and presentation. In PAL, we use the dynamic capabilities of SBUS to connect system components as appropriate to achieve particular functional goals. For instance, should a carer wish to visualise a patient's live heart rate data, their (presentation) application must be able to find and connect to the relevant ECG stream. Sensor streams might also need to be directed to an inference engine to detect a change in state and raise alerts; e.g. when Oscar has fainted and not moved. Some connections might be triggered by policy, e.g., automatically connecting an emergency service to a patient's data streams (e.g. location, vital-signs) in situations of a perceived emergency.

D. *Information governance*

Personal information is inherently sensitive. Therefore, we must protect data at all system levels. The patient must be able to describe the circumstances in which particular information is disclosed to particular parties. Such policy is context-sensitive, as restrictions will often vary according to circumstance.

For the scenarios of both modes to work, subsets of Oscar's records need to be available to a number of different groups. These may include families, friends and neighbours, local GP practices, hospitals, paramedics and ambulance crews. Members of each group should have access only to the appropriate parts of the records and only for the appropriate length of time. Health and welfare records may include many different items in many formats, e.g. emails, photos, sensor readings. To deal with this, PAL applies protection directly to the items of sensitive data. These protected items can then be safely stored and transported on any medium as the data remains safe wherever it is, such as on servers, wired or wireless connections or stored on a disk or USB devices. Access to the data is controlled according to policies defined by Oscar and enforced using policy based authorisation. This allows them to share his records while maintaining control over who can see them and under what circumstances.

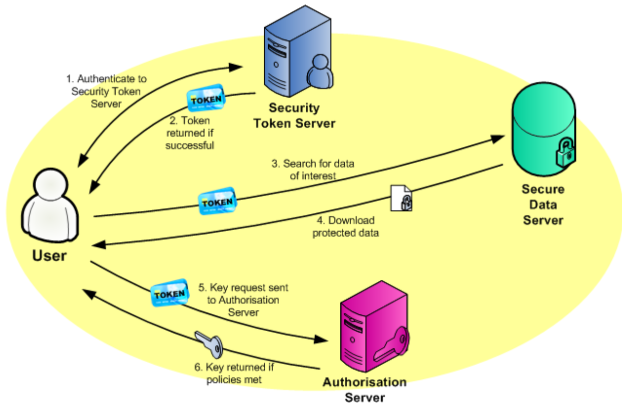

Figure 4. Token-based security mechanism

PAL users are identified within the system by using a token asserting the user's identity and attributes (e.g., a role). Tokens are issued by the Security Token Server (see Figure 4) once the user has been authenticated. When a user has a token they can obtain Oscar's data by connecting to a (live) data stream or by searching a Secure Data Server for archived data. The data is protected by encryption and the decryption key must be retrieved from the Authorisation Server. The Authorisation Server will check the credentials contained in the user's token and the information about the data object against the set of security policies. Access to the data will only be granted if the policy conditions are satisfied. A decryption key will then be returned to the user and they can decrypt the object and access the data. Since many parties are producing and consuming data within PAL, the system has been designed to operate within more than one administrative domain. Each domain can contain an instance of the components described above.

In addition to encryption, policy operates at runtime to govern information exchange. Each component maintains access control policy dictating the components with which it may interact. The SBUS middleware enables these privileges to be dynamically changed, which allows, for example, the relaxation restrictions in an emergency. Further, dynamically enforced policy enables granular, context-based control. For instance, policy might dictate that a homecare nurse is automatically disconnected from a patient's data streams when she physically leaves the patient's home. Both of these features are implemented by a dedicated, patient-centric policy engine that operates to trigger SBUS reconfigurations in response to changes in circumstance.

Patients undoubtedly have specific concerns, e.g., that their location should be hidden in normal circumstances, or that a GP cannot access their data unless explicitly authorised. 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 disclosure policy will be influenced by care providers, possibly through the use of policy templates that describe particular interactions. Any involved party is implicitly responsible for protecting health information, and thus must adhere to any specific patient preferences.

V. BRINGING THE PIECES TOGETHER

So far, within our project, we have focused on designing, developing and testing the main system components and functionality needed to realise the two main scenario modes (preventative and reactive) as well as the security and privacy framework described above. Our current work is to bring such pieces together within a combined prototype, following the integration framework. Here, we present details of our existing working system components and certain relevant results collected through our user-based experiments.

A. The Lifestyle Management System

Following the *preventative* part of our lead scenario, we developed a lifestyle management system that collects, processes and visualises various types of user information. For realising the system, we use various input sources for obtaining valuable information about a user's lifestyle. The recording scenarios had to reflect real lives as much as possible; therefore we considered a mixture of mobile and stationary sources, such as mobile phone, desktop, laptop, wearable physiological sensors and content servers. Through such sources, we are able to collect a large variety of information, as shown in Figure 5.

It is important to note that the types of information considered within a specific scenario highly depend on user preferences and constraints, the available sources within certain scenarios as well as on the relevance of information to such scenarios. A large amount of information is collected through the *AIRS (Android Remote Sensing)* platform [7], which benefited from our user-based experiments to extend towards incorporating more sensors as well as user-based annotations for mood and important events.

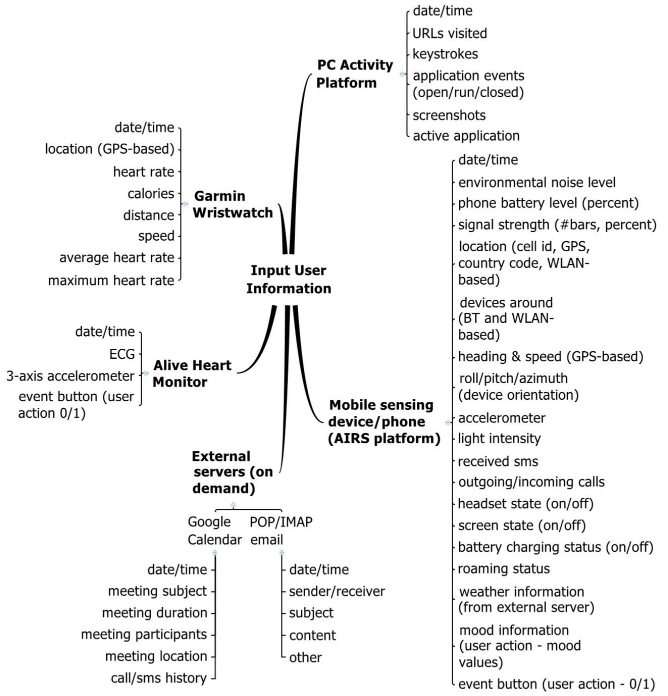

Figure 5. Information and sources.

The gathered data is secured and stored in a user-controlled database. Given that our solution is currently not embedded into a particular healthcare environment, the security token framework described in Section IV.D was not integrated.

Our system processes the gathered information along the context dimensions mentioned in Section IV.C for determining, for instance, activity, location, or environment information. This processing directly supports the visualisations towards the end user, both at detailed and abstract level. For this, the story-based approach presented in Section IV.C is realised (with Figure 3 showing an example story fragment). In addition, we allow end users to view detailed information at various levels, such as ambient noise level over time, crowdedness as determined by Bluetooth-enabled devices around and so on. The end users can have access to their data (either recorded or obtained on-demand) through a blog-based interface. This 'diary' style of visualisation lends itself to a natural form of end user interaction, allowing for the selection of particular days, archiving by month or amending gathered information with personal entries.

In addition to the preventative part of our lead use case, we also realised certain aspects of the *reactive* side, such as the information routing at middleware and infrastructure level (see Section IV.B). Hence, we utilise the information routing approaches, either through the SBUS middleware over an IP infrastructure or through directly routing on information within an information-centric networking test bed. We furthermore use our communication component of Section IV.B to ensure continuous connectivity in cases where individual networks become unavailable. This is important in particular for mobile scenarios. We also integrated this solution with the endpoint-oriented security component as shown in Section IV.D. With this, we can demonstrate the secure streaming of, say ECG data, from a patient to an emergency response center without connectivity disruption (if alternative connectivity is available, of course).

B. *Insights from User Experiments*

In order to better understand how to design and build systems such as proposed in this paper, we have performed various user-based experiments aimed at giving us more insight into: (1) what information is considered useful; (2) how information should be shown; (3) what interactions should be used. The experiments were both survey-based (online questionnaire) and hands-on. The online survey tested relations between people that reflect on causes for certain important aspects of their wellbeing, such as stress, and their attitude towards systems as described here. The participants were invited through various means (university staff and student mailing lists, Facebook and LinkedIn contacts, PAL project, etc.). As background, we consider participants' experience with self-monitoring systems and their attitude towards self-reflection more important than technological abilities or age.

The results included here are based on 38 participants, 7 with previous experience with self-monitoring, 29 without and 2 that did not specify. The participants were introduced to the preventative part of our demonstrator system from the beginning, through a stress-related usage scenario and a high-level system description. They were asked if they would find such a system useful for self-reflection. Encouragingly, 63% of all participants answered positively. We were particularly interested in how people we considered more self-reflective (i.e., people that think back often *and* think why something happened) feel about the system and 76% said they would find such system useful. Out of the people that answered positively, 74% had no previous experience with self-monitoring systems. Next, we asked if they would like a story-based representation created by the system based on their recorded data (we used a concept movie containing a day story built from events similar to Figure 3). 83.3% of the people that said they would find the system useful said they would like such interface. We also asked if they would like to customise their story. 95% of the ones that like both the system and the story answered positively, with most people wanting to customise characters, places, emotional states, and activities.

We further conducted a series of dedicated experiments with 6 unpaid volunteers (ages between 20 and 80) that aimed at finding out what people consider interesting in their daily activities For that, the users were asked to use digital annotation means we provided (a binary event button on the Alive heart monitor and, later, AIRS-provided phone *widgets:* small interface elements that allow interaction with an application without brining it to the foreground) in order to mark when something they considered interesting enough to be reflected within their lifestyle management system happened or was about to happen.

Essential in understanding how systems like ours should be created is that every user stated that they expect their interest in certain types of information to change within the various situations experienced during the day as well as over time and with increased system usage. Hence, it is expected that the system would provide end users with means to customise their stories over time, both in terms of what they include and in terms of how they look. This would allow them to better understand what could be of interest for them in ways that they have not thought of before. Part of this aspect of *configurability* also includes that each user wants to be in control of what is recorded. Also, being able to change their preferences over time allows a change of focus.

As a result of our experiments we have already added several new features to our system, such as the phone-based widgets used for annotating an event with user's own meaning and another widget that allows users to set their current mood. We consider such features essential for a system like ours, as the automatic inference processes would probably not be able to (always) capture the meaning of why an event was interesting (e.g., is it because the user is in a certain place, or in a certain situation, such as meeting). It is worth noting that users still highlighted the need for faster annotation through dedicated hardware in some scenarios.

VI. CONCLUSIONS AND FUTURE WORK

Pervasive healthcare systems face many real-life deployment challenges that are rarely or poorly addressed. This often leads to disruptions in end user experiences or imposing system constraints, leading to single device or single network solutions. We believe that providing solutions only for bespoke use cases will prevent the wide adoption of healthcare solutions in the future. It is here where our main contribution lies, namely to outline system-wide challenges for healthcare provisioning systems that are addressed through an open, systemic and holistic approach. The resulting PAL system operates over a variety of communication infrastructure solutions. It abstracts the specifics of the infrastructure through a unifying middleware, and is able to support the functional requirements of a wide-range of applications.

At the application-level, we have shown how the lifestyle assistant enables the users of the system to better understand the plethora of available data, by employing novel visualisation approaches that extend detailed graph-based information through story-telling approaches that are engaging and summarising. All of this is achieved within a framework for information governance that takes into account the private and sensitive nature of the information involved. Although our intention is not to build specific diagnostic tools, the lifestyle assistant supports patients in reflecting on their activity history and also assists more formal care processes; thus providing insights for healthcare systems in general.

In our future work, we will investigate the possibilities for conducting detailed studies relating to different aspects of the system, such as visualisation or the security framework. We will also continue to work with funding agencies to incorporate the lessons learned into future system requirements.

Acknowledgements

PAL is funded by the TSB and EPSRC under grant numbers TP/AN072C, TS/H000062/1 and TS/H000186/1.

References

[1] UK Department of Health, "Improving Chronic Disease Management," Report, 2004.

[2] D. Trossen and 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), pp. 1-8, August 2007.

[3] NHS Direct, http://www.nhsdirect.nhs.uk, [retrieved: Sept, 2012].

[4] D. Ingram, "Reconfigurable Middleware for High Availability Sensor Systems", Proc. ACM Distributed Event-Based Systems DEBS09, Nashville, TN, USA, pp. 20:1-20:11, July 2009.

[5] C. E. Perkins, "IP Mobility Support for IPv4, Revised," Internet Engineering Task Force (IETF), 2010.

[6] EU FP7 Project PURSUIT, http://www.fp7-pursuit.eu, [retrieved: Sept, 2012].

[7] D. Pavel, V. Callaghan and A. K. Dey, "From self-monitoring to self-understanding: Going beyond physiological sensing for supporting wellbeing". International Workshop on Pervasive Computing Paradigms for Mental Health (MindCare 2011), Pervasive Health 2011, Dublin, pp. 312 -315, May 2011.

[8] AIRS: An Android-based Sensing Platform: https://play.google.com/store/apps/details?id=com.airs, [retrieved: Sept, 2012].

[9] D. Trossen, D. Pavel, J. Singh, J. Bacon and K. M. Guild, "Information-centric Pervasive Healthcare Platforms", Proceedings of 4th International Conference on Pervasive Health, Munich, pp. 1-4, March 2010.

[10] IBM, "MQ Telemetry Transport (MQTT)", http://mqtt.org/, [retrieved: Sept, 2012].

[11] WristCare: http://www.istsec.fi/eng/Emikakoti.htm, [retrieved: Sept, 2012].

[12] SenseWear, http://sensewear.bodymedia.com/, [retrieved: Sept, 2012].

[13] Philips Lifeline Solutions, http://www.lifelinesys.com/content/home, [retrieved: Sept, 2012].

[14] Wellcore Emergency Response System solutions, at http://www.wellcore.com/, [retrieved: Sept, 2012].

[15] OBS, at http://www.obsmedical.com/products, [retrieved: Sept, 2012].

[16] Bosch Healthbuddy, http://www.bosch-telehealth.com/en/us/products/health_buddy/health_buddy.html, [retrieved: Sept, 2012]

[17] CardioNet Solutions, http://www.cardionet.com/patients_01.htm, [retrieved: Sept, 2012].

[18] Alive Technologies, "Alive Heart and Activity Monitor", http://www.alivetec.com/products.htm, [retrieved: Sept, 2012].

[19] MobiHealth, http://www.mobihealth.com, [retrieved: Sept, 2012].

[20] iFall, http://www.imedicalapps.com/2010/04/ifall-android-medical-app/, [retrieved: Sept, 2012].

[21] Sportstracker, http://www.sports-tracker.com/, [retrieved: Sept, 2012].

[22] Endomondo, http://www.endomondo.com, [retrieved: Sept, 2012].

[23] D. Trossen and G. Parisis, "Designing and Realizing an Information-Centric Internet," IEEE *Communications Magazine*, Special Issue on Information-centric Networking, pp. 60-67, July 2012.