Server-assisted Context-Dependent Pervasive Wellness Monitoring

Archan Misra, Ben Falchuk, Shoshana Loeb
Advanced Technology Services
Telcordia Technologies
Piscataway, NJ, USA

***Abstract*—Recent research on remote health monitoring solutions has focused largely on developing context-dependent, stream-processing capabilities on a personal mobile hub (typically, a cellphone) for energy-efficient transmission of data collected from a set of body-worn medical sensors. In this paper, we argue that commercial deployment of such pervasive wellness monitoring will require the extension of such 'context dependency' to the process of data collection (from the sensors to the mobile device) as well. In particular, the utilization of an individual's non-medical activity context, by the cellphone, in coordination with a backend server, is posited to be the key to supporting important objectives such as intermittent sensing and data security/privacy. The server-side platform is able to extract relevant context from a wide variety of personal or generic backend information streams; this context is then enriched by local context derived by the smartphone. We present an early outline of a middleware platform for supporting this objective.**

Keywords-smartphone, remote monitoring, privacy, context, energy-efficiency

I. Introduction

Remote medical monitoring is widely viewed as a compelling next-generation service opportunity by device manufacturers, telecom providers and application providers. Most monitoring architectures adopt the so-called 3-tier monitoring model[1,2] (illustrated in Figure 1), where a smart cellphone acts as a hub that collects data from body-worn sensors via a local personal area network (PAN) link (e.g., Bluetooth) and subsequently transmits the data, with or without additional processing, to a backend server. Such remote monitoring has been proposed for a variety of `patient' populations, such as those with chronic conditions (e.g., Alzheimer's or diabetes), the elderly (e.g., monitoring for gait irregularity) and even Army veterans (monitoring of post-traumatic stress symptoms). Existing remote monitoring research and prototypes can be broadly categorized into two extreme categories:

- Periodic monitoring typically involves the very infrequent sensing and transmission of some biomedical attribute (e.g., glucose or weight readings thrice a day). In these scenarios [3], the data rates are low and the cellphone simply collects the sensor readings and transmits them (without further processing) to the backend.
- Continuous monitoring approaches [4][5] focus on the continuous transmission of data from a variety of 'always-on' sensors. In these scenarios, the data rates can be quite high and event-processing middleware [6] is invoked to extract and transmit higher-layer significant events.

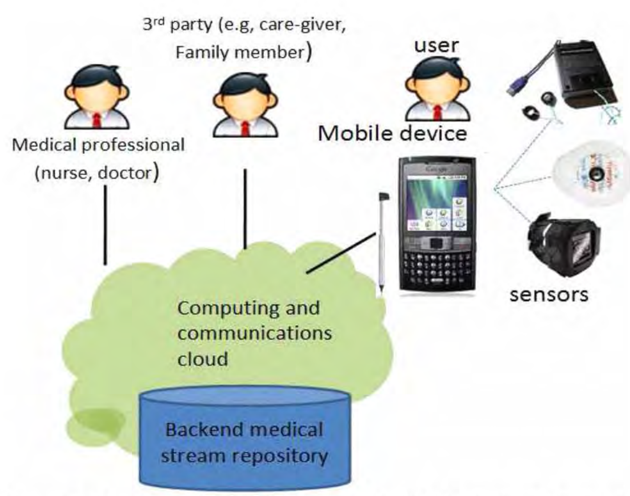

Figure 1. Basic 3-Tier Personal Pervasive Health Monitoring Infrastructure

We focus on a third, intermediate category of monitoring applications—those that are intermittent and episodic, yet require the transmission of continuous data streams during specific episodes. We denote this category as "**activity-triggered deep monitoring**" (*ATDM*) applications. These applications are characterized by a need to transmit potentially high-rate data streams, but only when specific 'contextual' patterns of the patient are either currently detected or expected in the future. This paper proposes that such ATDM applications will benefit from and require coordination of context monitoring both locally (on the cellphone) and global (by a backend server). In particular, we focus on demonstrating how *the actual process of data collection (i.e., the sensor→cellphone data transmission) can be intelligently adapted,* based on a collection of such local and global context. The key contributions of this paper are:

- We demonstrate how local and global context can together be used to coordinate the process of data collection for ATDM medical scenarios.

Digital Object Identifier: 10.4108/ICST.PERVASIVEHEALTH2009.5993
http://dx.doi.org/10.4108/ICST.PERVASIVEHEALTH2009.5993

- We next focus on a small set of technical aspects related to such adaptive monitoring, and identify areas requiring either technical innovations or standardized specifications.
- We introduce the basic functional components of our proposed ATDM middleware, and explain their role in adaptive context-aware data collection.

The rest of the paper is as follows. Section II discusses the characteristics of ATDM applications and the role of user context. Section III discusses several technical facets of such adaptive monitoring. Section IV then introduces the various components of our proposed middleware framework; finally, Section V concludes the paper.

II. Advances Needed for Commercially viable Adaptive Sensing

We first illustrate the characteristics of A*TDM* applications through a simple hypothetical example:

Consider an angioplasty patient who has been advised to preferably avoid strenuous exercise (that significantly stresses the heart) for a period of six months. The patient is using a Bluetooth-equipped heart-rate monitor (e.g., [10]) attached to her body supported by a networked context server (CS). Instead of wasting resources, an efficient data collection process would trigger data transmission only as certain activities are detected.

Now assume that the patient is ready to leave her office and travel to her next destination. As soon as the patient steps out of her office suite, the phone 'senses' that she is climbing up a few flights of stairs, and initiates the heart-rate monitoring to track the patient's cardiac state during this activity. Assume further that the CS is aware of her next appointment (from her corporate calendar entries) and infers that, because of local traffic congestion (obtained via real-time Web feeds), the journey will be stressful and will likely involve a brisk walk. The cellphone is instructed to keep the heart-rate sensor 'active' during this commuting duration.

The example illustrates the possibility that data collection may be intermittent and activity-driven. A key observation is that previous health monitoring scenarios have largely viewed the cellphone solely as a 'gateway' device. In general, *this viewpoint misses out on the increasing use of the cellphone as a 'personal activity coordination' device*, with a phone's scheduling capability (e.g., calendaring), as well as personal (e.g., food preferences) and generic context-awareness (e.g., GPS location), being combined for a variety of advanced applications in automotive telematics (e.g., personalized navigation services [8]), social networking (e.g., spontaneous social groups [9]) and other vertical domains. In addition, while localized context such as location and walking speed may be inferred by the mobile device, there is an increasing body of potentially valuable personalized or generic (e.g., traffic.com) context information that is stored in the network 'cloud'. Clearly, personal, Web 2.0 style services such as life-logging, location-logging (e.g., Yahoo! Fire Eagle), blogging, and social networking (e.g., Twitter, Tumblr) are allowing users to express and share their 'activities' (both physical and mental states) on many channels (e.g., niche wellness-related social networks such as DailyStrength.com and PsychCentral.com). Appropriate mining and reasoning (such as *Sentiment Analysis*) over this data may be used to refine or extract additional valuable context about a person's current (or future) activity levels (e.g., understanding that the scheduled "tennis" match infers high activity levels). This information may be used to both adapt the process of sensor data collection (the focus of this paper) and also perhaps initiate additional interactions with the individual (something not considered further in this paper).

Figure 2 thus illustrates a more extended vision of remote monitoring—in contrast to Figure 1, the monitoring infrastructure now includes a backend "Context Server" that is responsible for extracting, fusing, and reasoning over other related network-accessible context information.

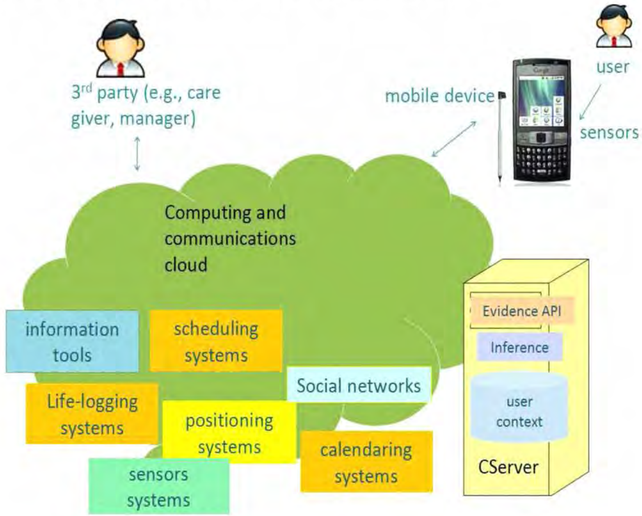

Figure 2. Extended Monitoring Infrastructure with Backend Context Server (CServer)

III. Technical Aspects of Commercially viable Adaptive Sensing

Most work on adaptive sensing and communication between a group of body-worn sensors and a personal hub have so far focused primarily on improving the communication efficiency (e.g., larger data rates, lower energy consumption). In contrast, our belief is that ATDM scenarios involve a few new forms of adaptation, which in turn are likely to require the creation of novel technologies or adoption of new standards by the cellular telephony industry.

A. Privacy and Security

Perhaps one of the most important uses of adaptation might lie in the enforcement of universal or personalized privacy and security preferences. An individual's medical streams are likely to be among the most privacy-sensitive pieces of electronic data—we believe that even isolated and rare incidents involving unauthorized access to such data may generate enough adverse public opinion to effectively stall the commercial adoption of remote monitoring technologies.

Digital Object Identifier: 10.4108/ICST.PERVASIVEHEALTH2009.5993
http://dx.doi.org/10.4108/ICST.PERVASIVEHEALTH2009.5993

Most prototypes of sensor→hub communication use a Bluetooth channel for communication; in general, it is known [11] that such channels may not provide the highest level of confidentiality. For improved confidentiality of PAN communications, one can impose additional data-layer encryption mechanisms on the data. However, there is an inherent tradeoff between the energy-efficiency and level of security in any communication channel. *Our view is that the phone must be capable of dynamically adjusting this interplay between energy vs. security, based on the 'context' of the patient.* For example, a lower-cost encryption technology may be adequate when the patient is in her private home; the phone may, however, switch to a more heavy-weight security mechanism when the patient is in a public space (e.g., a stadium or shopping mall). This is consistent with the trend of increasingly sophisticated sensing ability on the phone (e.g., [7][12] demonstrated how an onboard accelerometer could be used to distinguish between an individual's sitting, walking and running states) and the progressively diverse ways that a user's personal context can be inferred from the information she makes available on the Web (see Section IV)..

B. *Activity History*

ATDM applications occupy a unique niche, as the extent of medical data needed depends on the current and past activity history of the monitored patient. The interesting thing to note is that the actual implementation of such 'context determination' may be very different across different runtimes (due to differences, for example, in phone OS or set of attached sensors). Research prototypes so far have been developed as 'point-solutions', engineered to cope with the idiosyncrasies of each researcher's specific platform and sensor set.

To make monitoring applications 'device-independent' (an important step towards a commercially viable, large-scale monitoring practice), it is thus necessary for phone vendors/operators to define a common set of 'medically relevant activities' (e.g., via the development of ontologies) that abstract out the specifics of each platform and runtime environment, and then ensure runtime support for the recognition of such activity contexts. We expect that such standardization will involve extensions to existing 'activities' standards (e.g., the JSR75 PIM package or the CalConnect [13] consortium) that have so far focused on calendar-specific content.

C. *Programmable Activation of Sensors*

Central to the vision of ATDM applications is the ability to activate and de-activate individual medical sensors on demand (based on a user's activity context). In certain situations, such activation may be more deterministic (e.g., ten minutes of monitoring, once every hour). In addition to such basic control, advanced ATDM scenarios may require dynamic adjustment in a variety of other sensor-specific parameters. For example, to conserve energy, a mobile phone may instruct a sensor to adjust its sampling frequency, or perhaps, even its resolution level (an approach that might be more effective when the sensor aggregates multiple samples into a single data packet). Unfortunately, our experiences with most wireless medical sensors indicate a lack of such 'programmable' capability. (In many instances, sensors are simply turned on or off through manual 'button-pressing'.) To provide sufficient vendor-neutrality, the community will need to standardize specific adaptive behaviors on such sensors.

D. *Semantic Interpretation of Networked Information Sources*

Our proposed vision of remote monitoring requires significant extensions to our earlier work on context servers [8], so that servers are capable of using both structured and unstructured information streams to piece together a rich picture of a user's state (both current and future), ranging from physical to emotional. This is possible only in limited senses at the moment (see the variety of technologies supporting Web 3.0). Clearly, advances in information representation and analysis (such as blog sentiment analysis and natural language processing) will be needed to make the Context Server capable of the sorts of deep introspection and information correlation that are needed to extract usable and individualized context.

IV. CONTEXT-SENSITIVE MOBILE SENSING ARCHITECTURE

We have recently initiated promising work to build a mobile phone middleware that address the unique requirements of ATDM scenarios. Figure 3 shows the high-level components of our proposed middleware. The important components of the middleware on the phone include:

Personal Medical Policy Repository (PMPR) is the database where all the medical monitoring-related policies (e.g., security preferences, contexts of interest) for the individual are stored. We expect that the storage formats for each of these categories of information will be standardized to ensure vendor and service provider independence.

Activity and Context Trigger (ACT) is responsible for generating the specific activity contexts that may be relevant to specific monitoring applications. This component retrieves the specific 'context' states of interest from the Personal Medical Policy Repository (PMPR) and generates alerts when the corresponding context states are detected.

Privacy Enforcer (PE) is responsible for combining the privacy preferences in the PMPR with the context states detected by the ACT to determine the current set of 'policy/security' preferences.

Sensor Data Collection (SDCO) is responsible for the actual communication and reception of data streams from the individual sensors.

Dynamic Sensor Control (DSC) - controls the activation/deactivation of sensors, and the adaptation of their specified behavior (e.g., periodic activation every 1 hour and transmission of 8 bit sensor samples@200 Hz). Note that the DSC is logically separate from the SDCO (even though they use the same radio interface for communication)—the DSC focuses on adapting the sensor's 'sensing and communication' behavior, while the SDCO is responsible for the actual communication of medical data streams.

Digital Object Identifier: 10.4108/ICST.PERVASIVEHEALTH2009.5993
http://dx.doi.org/10.4108/ICST.PERVASIVEHEALTH2009.5993

Context-Dependent Event Processing Engine (CEPE) - aggregates and filters medical data streams, prior to transmission on the wide-area interface.

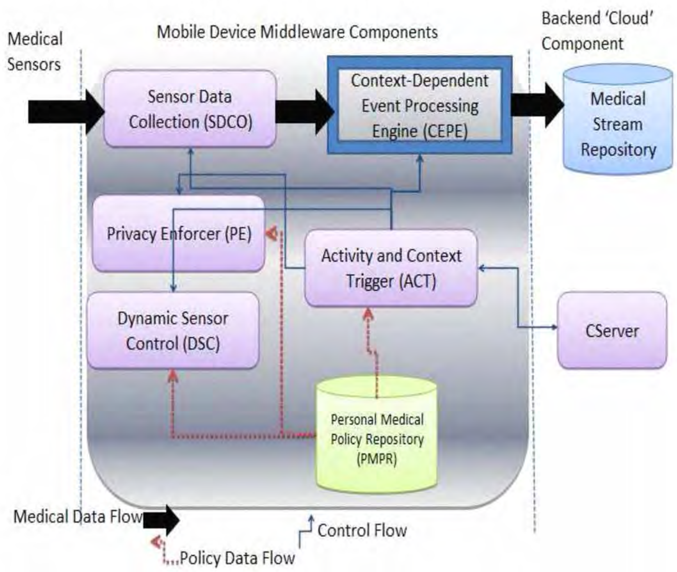

Figure 3. Functional Middleware Architecture

In addition to the client-side components, the figure also shows the backend *Context Server* (CServer). This Context Server is responsible for extracting relevant individualized context from information sources in the 'cloud'. The ACT component on the mobile device interacts with the CServer to obtain the most accurate 'context' of an individual's current or future predicted activity. Note that this separation of functionality promotes significant communication efficiency, by avoiding the need for the ACT on the phone to directly subscribe to (potentially large volumes of) data in the 'cloud'.

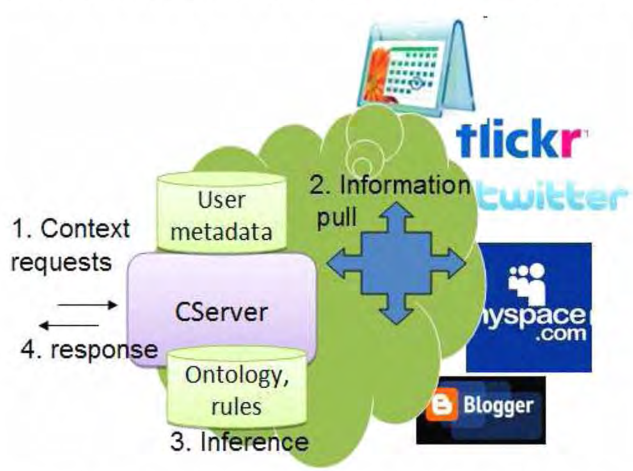

Figure 4. Fusion of Web data about a user in ordre to create a Projection

Figure 4 illustrates the CServer role with respect to the mobile device and the network. Depending on policy, the device middleware will occasionally ask the CServer for context information that helps it decide on stream management. CServer can be thought of as an external processing function that, by pulling and fusing information about a given user from (pre-configured) Web sources may provide additional insight into the user's current or imminent physical or emotional state. An ontology and heuristics help the CServer piece together a view of the user's state that will ultimately be helpful to the device middleware. As examples, the CServer might assert that the users privacy should be increased at the moment given all the conditions surrounding the user at the moment, or that the user is about to be in a state of high anxiety due to environmental conditions.

V. CONCLUSIONS

In this paper, we have introduced the concept of using rich 'user context', gleaned or inferred both locally by a mobile device and globally from 'cloud' information sources, to adapt the process of data collecting in pervasive health monitoring systems. Prior results for the Context Server (from a deployment [8] aimed at mobile operators) have shown that even subtle context, such as scheduled calendar events, can have significant operational benefits. We are currently working to build a prototype of our proposed middleware.

VI. ACKNOWLEDGEMENTS

We acknowledge colleagues Thimios Panagos, Russ Fischer, and George Collier for their input and discussions.

REFERENCES

[1] D. Husemann, C. Narayanaswami, and M. Nidd, "Personal mobile hub", In Proceedings of the Eighth International Symposium on Wearable Computers (ISWC'04), Washington, DC, USA, 2004.

[2] E. Lubrin, E. Lawrence, and K. F. Navarro, "Motecare: an adaptive smart ban health monitoring system", In Proceedings of the 24th IASTED Int'l Conf. on Biomedical engineering (BioMed 06), 2006.

[3] M. Blount, et al, "Remote health-care monitoring using personal care connect" *IBM Systems Journal*, 46(1):95–113, 2007.

[4] R. Jafari, F. Dabiri, P. Brisk and M. Sarrafzadeh, "Adaptive and fault tolerant medical vest for life-critical medical monitoring", Proc. 2005 ACM Symp. *on Applied computing*, pages 272–279, New York, 2005.

[5] E. Lubrin, E. Lawrence, and K. F. Navarro, "Motecare: an adaptive smart ban health monitoring system", In BioMed'06: Proc. of the 24th IASTED Int'l Conf. on Biomedical Eng., pages 60–67, Anaheim, 2006.

[6] I. Mohomed, A. Misra, M. Ebling, W. Jerome, "Context-aware and personalized event filtering for low-overhead continuous remote health monitoring", In Proc. of 9th IEEE Int'l Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM), June 2008.

[7] B.Falchuk, S.Loeb, K.Sinkar, A.Dutta, "Mobile Contextual Mashup Service for IMS", Proc. 2nd IEEE Int'l. Conf. on Internet Multimedia Services Architecture and Applications, Bangalore, Dec. 2008

[8] B.Falchuk, S.Loeb, T.Panagos, "A Deep-Context Personal Navigation System", Proc. ITS America 15th World Congress on Intelligent Transportation Systems, New York, 2008

[9] Y. Chang, H. Liu, L. Chou, Y. Chen and H. Shin, "A General Architecture of Mobile Social Network Services", In Proceedings of International Conference on Convergence Technology, November 2007.

[10] Nokia N79 and Polar Heart Rate Mointor, http://www.psfk.com/2009/01/ nokia-n79-active-tracks-heart-rate.html

[11] Y. Shaked and A. Wool, "Cracking the Bluetooth PIN", In Proc. 3rd USENIX/ACM Conf. Mobile Systems, Applications, and Services (MobiSys), pages 39-50, Seattle, WA, June 2005.

[12] D. Gafurov, E. Snekkenes, P. Bours, "Gait Authentication and Identification Using Wearable Accelerometer Sensor", IEEE Workshop on Automated Identification Advanced Technologies, June 2007

[13] The Calendaring and Scheduling Consortium, http://www.calconnect.org

Digital Object Identifier: 10.4108/ICST.PERVASIVEHEALTH2009.5993
http://dx.doi.org/10.4108/ICST.PERVASIVEHEALTH2009.5993