HARMONI: Motivation for a Health-oriented Adaptive Remote Monitoring Middleware

Iqbal Mohomed[1], Maria R Ebling[1], William Jerome[1], Archan Misra[1]

[1] TJ Watson Research Center, 19 Skyline Drive, Hawthorne, NY 10532
{iimohome, ebling, wfj, archan}@us.ibm.com

Abstract. This paper motivates the need for adaptive behavior in personal gateway devices (such as cell phones) used in remote health monitoring solutions. We put forth our vision of continuous monitoring of an individual's health, and consider the challenges presented by mobility in achieving this goal. We introduce the HARMONI middleware, which is designed to (i) enable the personal gateway device to provide useful monitoring even while disconnected from remote servers, (ii) learn patterns in sensor data to reduce resource requirements for relaying data, and (iii) tailor data relay based on anticipation of the user's environment and resource availability in the near future. The paper concludes with a status of our project.

Keywords: Remote Health Monitoring, Pervasive Healthcare, Adaptation, Mobile Devices

1 Introduction

Rapid advances in low-cost, wireless personal area networks (e.g., Bluetooth[1]) and wearable sensors (e.g., Human++ from IMEC[2]) foreshadow the emergence of remote health monitoring as a crucial tool for affordable future healthcare delivery and management. A limited number of personal health monitoring products are already available in the market [3][4][5][6]. While these products vary in sophistication and cost, ranging from basic systems that collect data wirelessly from biomedical devices and store it for later access to more advanced solutions that collect the data, distribute it to a managed care team and proactively trigger alerts for physicians, they are principally designed to support relatively infrequent, patient-triggered and non-adaptive monitoring.

The real power of such remote monitoring will, however, be harnessed once such monitoring becomes essentially *continuous*, allowing the biomedical parameters of both chronically ill or healthy patients to be collected, potentially over years or decades. It is expected that the ability to collect such a *detailed, fine-grained and continuous* medical history of an individual can revolutionize healthcare treatment, permitting advances such as proactive detection of emerging medical conditions, individually customized treatments, continuous safety monitoring of at-risk patients and detailed evaluation of patient response to medication. Such advances are important to limit the economic pressures imposed by rapidly aging populations in many societies of today.

The prevailing and most widely accepted architecture for such remote health care monitoring is illustrated in Figure 1, and involves the use of a personal gateway device (such as a PDA or smart phone) interposed between the body-worn biomedical sensors and the remote server acting as the repository of biomedical data streams. As the sensors typically have no wide-area networking interface due to energy and form-factor limitations, they transmit their sensed data to the gateway device via a short-range technology such as Bluetooth. This gateway device is assumed to possess wide-area wireless connectivity, and is responsible for relaying such transmitted data back to the remote server for either real-time or offline analysis. Such a three-tiered architecture has been proposed and prototyped in several research initiatives [7][8]. We believe that such a three-tiered architecture, especially one that leverages a ubiquitous consumer device such as the cell phone, is especially attractive, as it allows remote monitoring to become simply a value-added application utilizing an existing and universally available cellular infrastructure.

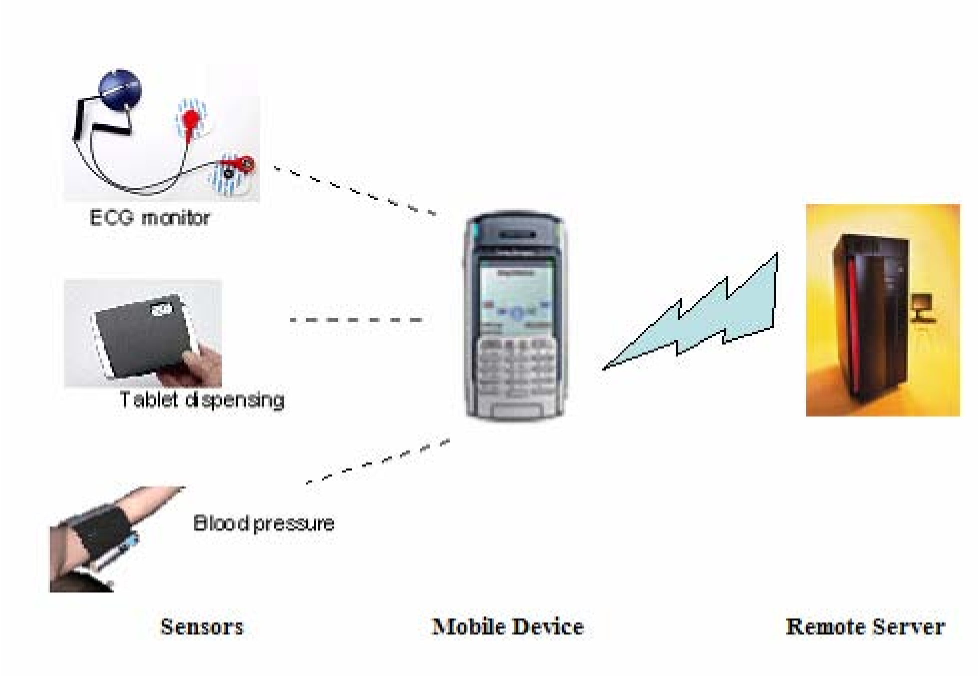

Fig. 1. Sensors transmit data to a personal gateway device (e.g. cell phone) via low power radio, which in turn relays the data to a remote server via a wide-area wireless interface.

However, all the existing proposals for such three-tiered monitoring hold the implicit belief that the pervasive gateway device (i.e., the cell phone) acts as a pure relay, in effect simply forwarding all the sensor readings unaltered to the backend server. If the data indicates a potential worsening in the patient's condition, the remote server can send an alert to the patient's mobile device or alert the patient's care team. If the mobile device experiences disconnection, it simply caches sensor data until a connection can be reestablished, which means that the patient is not being actively monitored during these periods.

In this paper, we discuss the challenges that arise when such a three-tiered architecture is used for long-term, continual, fine-grained biological monitoring. Next, we introduce HARMONI, a system for long-term monitoring that we are currently prototyping, in which the remote server and the mobile client software act in concert to adapt to the challenges presented by mobility. HARMONI is designed to enable the mobile client to provide useful monitoring even while disconnected from remote servers. HARMONI is also designed to make efficient use of resources such as bandwidth and energy. The system has the ability to learn patterns that occur in sensor data, and can use this technique to reduce the amount of data needed to be relayed from the mobile device to the remote server. HARMONI also tailors its data relay based on anticipation of the user's environment and resource availability in the near future. We conclude the paper with a summary of our current status and a discussion of how we hope to benefit from attending this workshop.

2. New Challenges with Long-Term, Continuous Data Relay

The vision of continuous long-term monitoring using a variety of rich wearable sensors imposes several challenges on the operation of the ubiquitous personal gateway (e.g., the cell phone).

2.1 Energy Limitations

One of the biggest drawbacks of current continuous monitoring architectures is the implicit requirement for the mobile device to transparently forward the entire data stream generated by the associated set of biomedical sensors. Such transparent monitoring is reasonable when the volume of generated medical data is low—for example, when the cell phone merely relays one or two readings daily that are generated when a monitored patient physically activates a glucose sensor. The increasing sophistication of sensors, however, points towards a *significantly more data-intensive future*—for example, prototypes of EKG, heart-rate and muscular activity sensors already exist that generate *quasi-continuous data*, with samples generated once every few seconds. In such situations, it is important to improve the energy-efficiency of the data collection and transmission process on both the mobile device and the individual sensors, to ensure that the operational lifetime of these battery-operated devices is not significantly curtailed. Our vision of practical healthcare monitoring assumes that the monitoring application piggybacks on the cell phone which is primarily used for conventional communications/entertainment applications such as voice and instant messaging. Accordingly, the focus on reducing the energy footprint of the application is especially crucial, to ensure that the continuous biomedical monitoring occurs almost transparently to the user, without any perceptible impact on the functioning of the phone.

2.2 Coping with Disconnection

Phones and PDAs occasionally encounter prolonged periods of disconnection from the cellular network. When the mobile device loses connectivity, it is unable to relay sensor data to the remote server. However, the device can store any sensor data received in the interim, and forward it to the remote server upon reestablishment of connectivity. Of course, under conditions of prolonged disconnection, the mobile device may need to partially discard the locally stored data streams due to insufficient availability of local storage. Additionally, loss of connectivity also limits the ability of the backend infrastructure to control the behavior of the phone or trigger other local adaptations in response to sensed data (such as triggering alerts on the mobile device if the biomedical readings lie outside specified thresholds).

2.3 Exploiting Multiple Network Links

Cell phones and PDAs are increasingly equipped with multiple network interfaces (e.g., a combination of 802.11, Bluetooth and cellular radios). Mirroring this trend, cellular operators are also increasing the range of supported access technologies (for example, combining wide-area GPRS service with 802.11 hotspot services, or enabling cellular phones to function as cordless devices using the fixed-line infrastructure when located at home). While this provides the device and user with increased flexibility, the monitoring infrastructure must be enhanced to exploit this access heterogeneity. Different interfaces have varying energy consumption and bandwidth costs associated with their use. A naïve approach that transmits data continuously over a pre-designated interface is unable to take advantage of a more efficient network interface that may be available in the near future.

2.4 Economics of Data Transmission

Any application based on wireless transmission of even moderate amounts of data must deal with the reality that usage of wireless spectrum will always incur some monetary cost. Many service plans for cellular data services (e.g., most schemes in the US) charge based on the amount of data that is transmitted over the link. One of the challenges with continuous, remote healthcare monitoring in a multi-technology, multi-provider environment is that the economic cost is not necessarily constant or predictable. For example, as the monitored user roams across different countries or regions, the cost associated with data transmission may fluctuate significantly (e.g., international “roaming” charges on data traffic are often prohibitively expensive). As an illustration of the economic impact of the choice of communication interface, note that 802.11-based data hotspots have experienced significant market penetration globally. While most 802.11 access models offer flat-rate data access and are thus more attractive for intensive data transmission, this advantage must be balanced against the reality that cellular coverage will remain more universal than WiFi “hotspots” for the foreseeable future.

3 HARMONI Innovations

The HARMONI system is being designed and built to achieve long-term, continuous data relay. To achieve this end, HARMONI has three key features. First, it can provide useful monitoring even while disconnected from the remote server. Second, it is able to learn commonly recurring patterns in the sensor readings recorded for a particular individual. When a known pattern is identified on the mobile device, it may be sufficient to indicate to the remote server that the pattern was observed rather than transmitting all of the collected sensor data. Third, the system can anticipate events that affect resources, and use its expectations to tailor data transmissions to the remote server. Next, we describe these three features in detail.

3.1 Useful Monitoring While Disconnected

HARMONI is being designed to provide more than just data-relaying functionality. The system can handle disconnection, not just through effective policy-based management of locally stored data awaiting upload, but also using intelligence on the phone to perform a basic level of local decision-making.

Broadly speaking, the client is responsible for performing data collection, relatively simple recognition of pre-specified spatio-temporal patterns in this sensed data and triggering of appropriate actions based on matches to specific medical event patterns. The server on the other hand is responsible for aggregating external context, and combining this with individualized analysis of past medical history to determine appropriate rules (composed of patterns in data readings and actions to perform when the pattern is encountered) for the client device, and for dynamically informing the client device of changes in such event-driven rules.

When disconnection occurs, the client device can continue to provide monitoring functionality, as it is operating on rules that have been relayed previously. Of course, the sophistication of monitoring under disconnection will not be equivalent to that under communication with the remote server, as the server has far more resources, and access to richer information sources than the mobile client.

3.2 Learning Recurring Patterns in Sensor Readings

Despite inherent variations, sensor readings typically conform to a “normal” pattern unless there is a marked change in the user’s overall health condition. The pattern of readings comprising the norm will almost certainly differ across individuals, and will be closely correlated to the user’s context. For example, a user named Alice may typically have a heart rate between 75-80bpm (beats per minute), except when she goes to the gym, at which time her heart rate is between 125-130bpm.

If we can learn commonly recurring patterns within the user’s sensor readings, we can achieve significant reductions in the amount of bandwidth consumed on the wireless link between the mobile device and the remote server, and also save on energy costs of using the radio link. Instead of transmitting raw or batched sensor readings, the mobile device can simply transmit an identifier for the pattern of readings that was observed.

Even though relaying patterns in lieu of actual observations can lead to resource savings, it is important to note that this scheme decreases the granularity of data that is ultimately collected on the remote server. Due to medical reasons, it may be necessary to transmit observations at the best possible granularity for some values of sensor readings or for certain contexts. Generally speaking, different situations may permit varying granularity in the collected data, and HARMONI takes this into consideration.

3.3 Anticipating Events that Affect Resources

Over the course of a day, there are certain events that have a significant impact on the resource levels of a mobile device. Two such events are recharging batteries and synchronizing data. We expect that users perform both activities with some regularity. For example, a user may typically recharge her cell phone at night. However, if she forgets to do this, she is very likely to charge her phone the next day at the office.

There may also be some regularity in the network connectivity available to a user’s device. For example, a user may work indoors and get no WIFI or cellular coverage in her workplace. However, at lunch time, the user visits the cafeteria of the building and has free WIFI connectivity for 30 minutes. Finally, when the user leaves the office and goes home, she may have good cellular connectivity.

HARMONI takes advantage of the regularity in battery recharge and data synchronization events, as well as historical patterns in the availability of various network interfaces. We have designed an anticipation mechanism that uses history, as well as current context information and resource levels on the mobile device, to decide whether data should be transmitted to the remote server immediately, or whether transmission should be deferred. For example, the system can choose to defer transmission of data when it believes a more efficient mechanism will be available in the near future.

It is important to note that medical reasons may necessitate that some values of sensor readings may require immediate transmission to the remote server, and this is taken into account by HARMONI's anticipation mechanism.

4 Status

At present, our group has a working implementation of a remote health monitoring system in which the mobile device (implemented on a Sony Ericcson P900/P910 cell phone) acquires data from a variety of health sensors via Bluetooth, and performs periodic transmission to a remote server via CDMA1X. The sensors we are using include an A&D blood pressure cuff (Model UA-767 Plus), an A&D weight scale (Model UC-321BT) and a Nonin Pulse Oximeter (Bluetooth-enabled Avant 4100 Wrist-Worn Module). Furthermore, two live pilots of this system are currently underway. However, the working implementation is designed primarily for user-triggered sensor monitoring and relatively infrequent data collection—it does not provide the smart, adaptive control and local intelligence envisaged in this work.

We are currently developing a prototype HARMONI system, which uses a more advanced Nokia pervasive mobile device and incorporates additional wearable sensors that provide a richer set of quasi-continuous biomedical data. As part of this development, we are designing algorithms to drive the disconnected operation, pattern learning, context anticipation and intelligent data compression components. After completing the prototype development in the next couple of months, we plan to evaluate it with actual users.

In attending the workshop, we hope to get feedback on the usefulness of our approach within a healthcare setting. We will also get the opportunity to showcase our research to date, and interact with other groups working to apply pervasive technologies to the domain of healthcare.

References

1. The Official Bluetooth Web site, http://www.bluetooth.com/bluetooth/
2. Interuniversity Microelectronics Centre (IMEC), http://www.imec.be
3. Honeywell HomMed, http://www.hommed.com
4. Health Hero Network, http://www.healthhero.com
5. American Telecare, http://www.americantelecare.com
6. AMD Telemedicine, http://www.amdtelemedicine.com
7. Dirk Husemann; Chandrasekhar Narayanaswami; Michael Nidd, "Personal Mobile Hub" 8th International Symposium on Wearable Computers (ISWC 2004), 31 October - 3 November 2004, Arlington, VA, USA
8. R. Chakravorty, "MobiCare: A Programmable Service Architecture for Mobile Medical Care", Proceedings of UbiCare 2006: The First Workshop on Ubiquitous and Pervasive Healthcare, March 2006

Author's Biography

Iqbal Mohomed *(iimohome@us.ibm.com/iq@cs.toronto.edu), IBM Research Division, Thomas J. Watson Research Center.* Mr. Mohomed is a PhD candidate at the Computer Science Department in the University of Toronto with an interest in pervasive computing, mobile devices and the application of computing technology to problems in healthcare. Presently, he is an IBM Research Intern at the TJ Watson Research Center. Mr. Mohomed has lectured several courses at the University of Toronto including Information Systems Analysis and Design, Software Tools and Systems Programming, and Discrete-Event Simulation and Modeling. In 2000, Mr. Mohomed received a B.Sc. Hons. from the University of Toronto with a specialist in Software Engineering, major in Economics and a minor in Mathematics. In 2001, Mr. Mohomed received a Graduate Certificate in Intelligent Information Systems from the Illinois Institute of Technology. That year, he was also a visiting student studying Game Theory at Northwestern University. In 2004, he completed an M.Sc. in Computer Science from the University of Toronto. Mr. Mohomed has designed and developed large software systems at Nortel Networks and Motorola. He has also been involved in developing application software for CaseWare International and Pervasive Software.

Maria R. Ebling *(ebling@us.ibm.com), IBM Research Division, Thomas J. Watson Research Center.* Dr. Ebling is a Research Staff Member and manager of the Privacy-enabled Context Technologies department at the Thomas J. Watson Research Center. She received a B. S. degree in mathematics with a computer science option from Harvey Mudd College, and M. S. and Ph. D. degrees in computer science from Carnegie Mellon University. She subsequently joined IBM Research, where she has worked in the area of context-aware computing, most recently in applying context-aware computing in the area of healthcare. In 2003 she received an IBM Technical Group Award for her work on intelligent notification. Dr. Ebling is a member of the Institute of Electrical and Electronics Engineers Computer Society and the Association for Computing Machinery.

William F. Jerome *(wfj@us.ibm.com), IBM Research Division, Thomas J. Watson Research Center.* Mr. Jerome is a Senior Software Engineer in the Privacy-enabled Context Technologies Department at the Thomas J. Watson Research Center. He received a B.S. degree in computer science from Brooklyn College and an M.S. degree in computer science from Columbia University. His recent research projects have been in presence awareness, healthcare and intelligent notification. Mr. Jerome is a member of the Association for Computing Machinery.

Archan Misra *(archan@us.ibm.com), IBM Research Division, Thomas J. Watson Research Center.* Dr. Misra is a Research Staff Member at the IBM TJ Watson Research Center, Hawthorne, NY, where he works on infrastructural components for pervasive networking and middleware for presence-based, context-aware computing. His current research projects include middleware for SIP-based "rich presence" in converged networks, high-performance wireless mesh networks and sensor data management techniques for healthcare and remote monitoring scenarios. Archan received his Ph.D. in Electrical and Computer Engineering from the University of Maryland at College Park in May, 2000, and his B.Tech in Electronics and Communication Engineering from IIT Kharagpur, India in July 1993. Prior to joining IBM in 2001, Archan worked at Telcordia Technologies on architectures for QoS in the Internet and adaptive IP-based mobility management algorithms. He has published extensively in the areas of wireless networking and pervasive computing and is a co-author on papers that received the Best Paper awards in ACM WOWMOM 2002 and IEEE MILCOM 2001. He is presently on the editorial board of the IEEE Wireless Communications Magazine.