Proceedings of AICA 2011,Torino, November 15-17, 2011.

Real-Time and Energy Issues in Mobile Health Monitoring Systems

Giorgio Buttazzo, Mauro Marinoni, Gianluca Franchino
Scuola Superiore Sant'Anna
Via Moruzzi 1, 56124 Pisa
{g.buttazzo, m.marinoni, g.franchino}@sssup.it

Health care is one of the areas in our society in which embedded computing systems can have a tremendous impact in reducing costs, improve efficiency, speed up the diagnosis of diseases, and detect critical health conditions. This paper presents an advanced mobile monitoring systems for patients affected by heart failure, focusing on the software methodologies adopted for handling real-time and energy constraints.

1 Introduction

Mobile health monitoring combines medical devices for health care with mobile computing technologies for improving the communication among patients, physicians, and other health care centers.

Recent technological advances in sensors, integrated circuits, and wireless communications have enabled the development of low-cost, miniature bio-medical devices that can be integrated into wireless body area networks (WBANs) for mobile health monitoring. Such networks are capable of sensing, processing, and communicating a set of vital signals and physiological data directly to a medical service center, where they can be stored into a personal file, which can be remotely accessed by the doctor whenever needed.

The advantages of such systems are significant, both from the patient and the society. From the patient point of view, this technology enables the delivery of accurate medical information from anywhere at any time, providing a snapshot of the personal health status to the doctor, without requiring hospitalization. Hence, the patient can continue to carry out his activities in his home, saving a lot of time and also gaining a psychological benefits. Moreover, the acquired data are stored in a personal file and collected to build up a complete and detailed history of the patient, useful for anamnesis, diagnosis, and prognosis. Finally, this technology helps people participating more closely in their own health care, especially for those that are reluctant to go to the hospital or to their doctor for routine medical checkups.

From a social perspective, mobile health care allows saving an enormous amount of money, by avoiding the costs of pre-hospitalization of those patients that need to be monitored but are not critical. Avoiding the pre-hospitalization of non critical patients has also the benefit of a more efficient usage of the resources, avoiding crowding and long queues in health centers and hospital departments.

A portable health monitoring systems must satisfy several requirements in order to be really useful and accepted by the patient. First of all, it must be small, lightweight, and it must not consume much energy, having an autonomy of at least 24 hours. Such features impose tight constraints on the system architecture, especially on the processing unit, the memory, and the wireless communication device. Given the strong limitations on the processing resources available onboard, the operating system must adopt suitable algorithms for efficient as well as predictable resource management.

Moreover, the system has to carry out several concurrent activities in real-time, most of which require to be periodically executed (e.g., sensory acquisition, filtering, digital signal processing, data integration, storage, and communication). As a consequence, the kernel must handle periodic tasks with timing constraints and must guarantee bounded response times for all time critical processes.

Finally, these systems are typically required to monitor the patient continuously from one up to 3 days, hence battery lifetime is a crucial constraint to be guaranteed. Unfortunately, real-time and energy constraints are contrasting objectives in terms of resources allocation. In fact, to achieve short response times the system should operate at high speeds and have all devices active to reduce access delays, but this generates high power consumption; whereas, in order to last long, the system should to operate at low speed and turn off devices whenever not used, but this generates long latencies. To reduce power consumption while guaranteeing real-time constraints, special energy-aware policies have be used for tasks and message scheduling.

In this paper, we present an advanced mobile monitoring systems for patients affected by heart failure, focusing on the software methodologies adopted for handling real-time and energy constraints.

1.1 Related work

In the last ten years, health care triggered several research projects related to mobile monitoring systems.

At Harvard University, researchers have developed a wireless monitoring systems, called Code-Blue [Mal04], for medical applications, including pre-hospital and in-hospital emergency care. The sensors include portable 2-lead ECG, pulse oximeter, wearable Pluto mote with built-in accelerometer, and a module with accelerometer, gyroscope, and electromyogram sensor for stroke patient monitoring.

At the University of Aarhus, a research team developed and deployed a monitoring infrastructure for hospitals, including a location tracking system and context-aware applications running on interactive displays and mobile phones [Bar06].

MobiHealth [Mobi] aimed at the development of a mobile monitoring system for healthcare, integrating sensors in a body area network (BAN) customized to the individual patient needs. Physical measurements, such as blood pressure or ECG are measured and transmitted wirelessly from the BAN to the hospital or health call center.

HEARTFAID [Heart] focused on defining an efficient health care delivery organization and management models for improving the processes of diagnosis, prognosis and therapy provision in the heart failure domain. The core of the platform is a clinical decision support system, designed by integrating knowledge representation techniques, hybrid reasoning methods, and advanced tools for the analysis of diagnostic data. In this context, a research team of the project [Col07, Col08, Chi10] investigated the use of semantic web technologies for implementing a real clinical scenario, covering the clinical course of a heart failure patient.

The ASCOLTA project [Asco] is focused on the development of an advanced monitoring systems for heart failure, integrating several physiological parameters like ECG, breath signal, oxygen saturation, blood pressure, inertial data, and patient weight, for early detection of critical conditions.

Stankovic et al. [Sta05] discuss the benefits and the crucial requirements for a pervasive monitoring infrastructure dedicated to assisted living, identifying the main research issues in such a new application domain.

Virone et al. [Vir06, Vir07] proposed an information system for assisted living based on a residential wireless sensor network that detects the presence of patients in specific locations and estimates a set of events of interest.

Wu et al. [Wu07] present the design and implementation of a wireless mobile monitoring system based on Bluetooth technology. Multiple physiological parameters can be monitored in real-time and the users can connect to the medical service center through digital family equipments, such as digital televisions and cell phones.

Wood et al. [Woo08] present a context-aware wireless sensor networks for assisted-living and residential monitoring, integrating environmental, physiological, and activity sensors in a scalable heterogeneous architecture.

Trappey et al. [Tra09] developed a mobile intelligent medical system to support mobile nursing applications and clinical decision support. The system includes an RFID-based module for patient data collection and a Java-based expert system for issuing warnings and diagnostic messages.

Ren et al. [Ren10] present several techniques that can be used to monitor patients effectively and enhance the functionality of telemedicine systems, and discuss how secure strategies can be used to improve the security of mobile healthcare.

Santra [San10] discussed secure strategies that can impede the attacks faced by wireless communications in mobile healthcare systems.

Shahriyar et al. [Sha09] present a mobile health monitoring system that uses a wearable wireless body area network for collecting data from patients and predicts patient's health status. The patients can access their health information by their mobile devices.

2 System architecture

The system considered in this paper is shown in Figure 1 and consists of an embedded platform worn by the patient to monitor his physiological data through a set of sensors. In its full configuration, the platform can include the following components:

- a microcontroller unit (MCU) in charge of acquiring, processing, communicating, and storing data from a set of biomedical sensors;
- a radio transceiver used to connect the system with a remote server or other devices, such as PCs or portable devices. Different wireless standard protocols are supported, as ZigBee, Bluetooth, and WiFi;
- one or more storage components, such as flash memories, SD cards, or USB storage devices;
- a set of biomedical sensors for monitoring the electrocardiografic signal (ECG), oxygen saturation (SPO_2), heart and breath rate;
- a 3-axis accelerometer to detect possible dangerous situations (e.g., falls) and reconstruct the patient activities for a correlation with the ECG trace.
- A user interface used to signal the system status and/or to control the device.

The system has been designed to be modular, in the sense that, depending on the specific application, one or more components can be excluded from the platform to minimize the weight and the energy consumption.. For instance, when the system is used to monitor people affected by heart failure, the platform includes ECG, SPO2, and breath rate sensors, a SD card to store the acquired data, and a WiFi connection to send the patient status to a remote health-care center. The user interface, kept as simple as possible, includes a switch to turn the device on and off, and a multicolor led and a buzzer to verify the proper operational modes. The system is powered by a rechargeable lithium polymer battery.

3 Real-Time issues

The software of a complex embedded device, such as that described in this work, must be able to handle several concurrent activities with different execution rates. For instance, the ECG signal is typically acquired with a sampling period of 2-4 milliseconds, while the breath rate sensor can be sampled with a period of 50 milliseconds. Other activities with different execution times and activation rates are related to data communication and data storage.

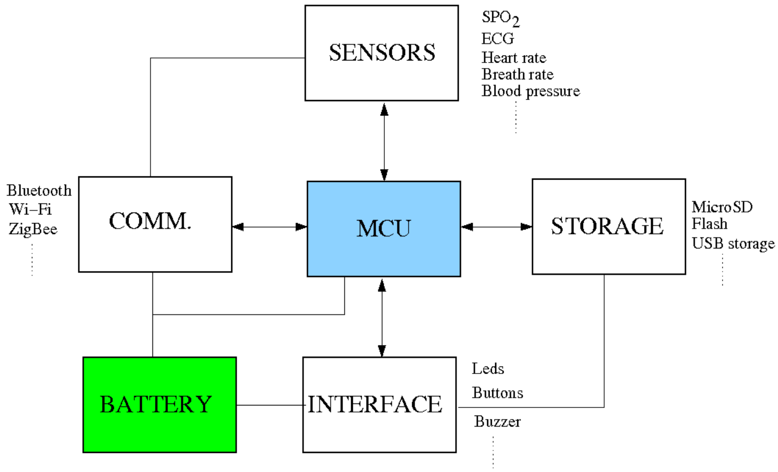

Fig.1 – System architecture

The data storage function is necessary to preserve patient data when the access point is out of reach(e.g. when the patient is not at home, or the communication is temporarily unavailable). The communication between the device and the remote health-care center is handled by a TCP/IP connection. To ensure a timely execution of such different activities and perform an off-line guarantee of the timing properties of the application, the software has to be supported by a real-time operating system. The software developed for the platform presented in this paper has been developed under the ERIKA Enterprise real-time kernel [Erika], which allows achieving high predictable timing behavior with a very small runtime overhead and memory footprint. The kernel is briefly described in the following section.

3.1.1 ERIKA Enterprise

ERIKA Enterprise is an open-source (GPL2 with linking exception) real-time operating system, implementing an Application Programming Interface compatible with the OSEK/VDX standard for automotive embedded controllers. ERIKA includes highly predictable real-time kernel mechanisms and uses innovative programming features to support time sensitive applications on a wide range of microcontrollers and multi-core platforms. In particular, ERIKA supports periodic and aperiodic task scheduling according to fixed and dynamic priorities; interrupt handling for urgent peripherals operation (interrupts always preempt task execution); and time bounded resource sharing through the Immediate Priority Ceiling protocol [Sha90] [But11].

4 Energy management

4.1 General issues

Power management with energy efficiency considerations is not only useful for mobile devices for prolonging their battery duration, but it is also helpful for server systems for the reduction of power bills. Dynamic power due to switching activities and static power due to the leakage current are two major sources of energy consumption in CMOS circuits.

For micrometer-scale semiconductor technology, the dynamic power dominates the power consumption of a processor. However, for technology in the deep sub-micron (DSM) domain, the leakage power consumption is comparable to or even more than the dynamic power dissipation. Reducing the voltage decreases the dynamic power consumption, however it also reduces the maximum operating frequency.

The overall energy consumption of a computing system also depends on other components. Martin et al. [Mar01] derived the following model to describe the power consumption as a function of the processor speed, defined as a normalized frequency ($s = f/f_{max}$):

$$P(s) = K_3 s^3 + K_2 s^2 + K_1 s + K_0$$

When the dynamic power consumption dominates the other components (due to switching and leakage), the K_3 coefficient is greater than the others and, power consumption can be approximated as $P(s) = K\, s^3$.

Two major classes of power-aware algorithms can be distinguished based on the kind of power they try to reduce. Dynamic Voltage Scaling (DVS) techniques reduce dynamic power by decreasing the supply voltage (and consequently the clock frequency) of the system, whereas Dynamic Peripheral Management (DPM) approaches achieve energy saving by exploiting operational states with reduced energy consumption (e.g., sleep or idle) whenever possible.

To reduce the runtime overhead due to energy management, some algorithms are applied off-line to compute the best energy setup as a function of the task set parameters (e.g., worst-case execution times and periods). The disadvantage of off-line approaches, however, is that they are based on worst-case assumptions and cannot take advantage of situations occurring at run time (such as early completions). For this reason, other methods are applied online to re-compute the parameters after each job termination.

When using DVS techniques in the presence of resource constraints (e.g., under mutually exclusive resources or non preemptive regions), the system can experience scheduling anomalies in which a task could even increase its response time when executed at a higher speed [But06]. Such problems prevent managing the performance of a real-time application as a function of the processor speed.

4.2 Proposed approach

In embedded systems characterized by real-time and energy constraints, selecting the most appropriate energy management policy is not easy, because the result heavily depends on the platform characteristics (e.g., energy modes and profiles of the devices, frequency range and power states of the CPU) and the application constraints (e.g., task deadlines, sensors acquisition delays, communication bandwidth, etc).

A preliminary study carried out on the presented platform showed that to achieve the best results, both DVS and DPM techniques have to be integrated to fit the requirements with the minimum energy cost. The combination of DVFS and DPM is done by forcing a CPU sleep interval followed by an active interval executed at a fixed frequency, selected to minimize the energy (per unit of computation) between the current and the next invocation of the analysis [Mar11]. The approach has been implemented in the Erika as a policy in the Power Manager of the kernel, which is described in the next section.

5 Kernel support

The energy saving module proposed in this section is part of the kernel and interacts with the scheduler, the hardware devices, and the application. While the scheduler selects the next task to execute, the Power Manager chooses an appropriate running configuration (i.e., speed and voltage). It uses a modular approach that allows the user to select a policy customized for a specific device providing a uniform interface.

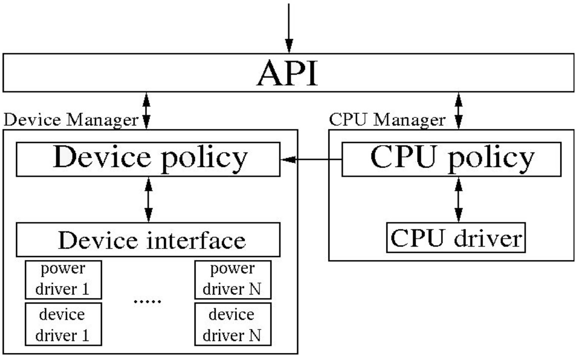

Fig.2 – Architecture of the power management module

The Power Manager consists of three hierarchically organized modules: the Application Programming Interface (API), the CPU Manager and the Devices Manager. Such a modular implementation allows the programmer to easily remove sub-components when not needed by the application, so helping to reduce the footprint.

The API module implements the interface defined for the interaction with the kernel and the applications. The CPU Manager is responsible for the power management of the CPU. Using a set of special callback functions the kernel informs the module about scheduling events.

The CPU policy sub-module implements the energy saving policies, which typically select the best speed to meet the applications constraints, while satisfying a given set of performance requirements. The CPU driver is in charge of setting the CPU parameters.

The Devices Manager handles internal and external peripherals. Inside it, the Device policy sub-module contains all the device policies. For each of them, two stacked components, Power driver and Device driver, abstract the device behavior using a discrete set of states.

6 Performance results

This section presents some experimental results that show that integrating DVFS and DPM approaches leads to lower power consumption in different types of platforms. The experiment compares a pure DPM approach, a pure DVFS approach, and the mixed one proposed by Marinoni et al. [Mar11] on a system that has to remain active during communication slots.

The comparison is done on two different types of architectures: in the first one (denoted as Fully-DPM model) the leakage power consumption is dominant; whereas in the second one (denoted as Fully-DVFS model) the main source of energy consumption is due to dynamic power.

In the experiment, the task set utilization is varied between 0.3 and 0.9 to evaluate the effect of the workload on different scenarios. The results are reported in Figure 3 and show that the mixed approach always produces the best performance in terms of energy reduction and, in the worst case, behaves as good as the other approaches.

Notice that, even in the most favorable condition (left graph), the fully-DPM approach presents poor performance working with energy constraints that prevent from switching the system off in some particular intervals.

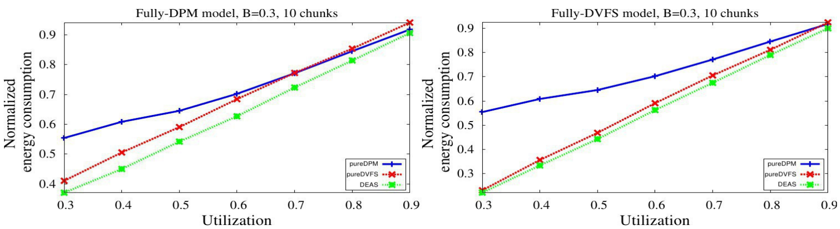

Fig.3 – Experimental results on power management techniques

7 Conclusions

This paper presented an integrated approach for addressing energy consumption and real-time constraints in a small embedded monitoring system for medical applications. Experiments results demonstrate the effectiveness of the method for different types of architectures.

8 Acknowledgments

This work has been carried out within the ASCOLTA project, funded by Tuscany Region under the grant 2009DUA-864500467-1.

9 References

[Asco] "ASCOLTA: Assistenza domiciliare allo SCOmpenso cardiaco attraverso Le Tecniche Avanzate di comunicazione digitale", Research Project by Tuscany Region, URL: http://www.ascolta.org/.

[Bar06] J. E. Bardram et al., "Experiences from Real-World Deployment of Context-Aware Technologies in A Hospital Environment," Proceedings of Ubicomp 2006, LNCS 4206, pp. 369–386, 2006.

[But06] G. Buttazzo, "Achieving Scalability in Real-Time Systems", IEEE Computer, 39(5):54-59, May 2006.

[But11] G. Buttazzo, "HARD REAL-TIME COMPUTING SYSTEMS: Predictable Scheduling Algorithms and Applications", Third Edition, Springer, 2011.

[Chi10] F. Chiarugi, S. Colantonio, D. Emmanoulidou, M. Martinelli, D. Moroni, O. Salvetti, "Decision Support in Heart Failure Through Processing of Electro and Echocardiograms", Artificial Intellingence in Medicine, Elsevier, Vol. 50, No 2, October 2010.

[Col07] S. Colantonio, M. Martinelli, D. Moroni, O. Salvetti, F. Perticone, A. Sciacqua, D. Conforti, A. Gualtieri, "An Approach to Decision Support in Heart Failure", 4th Italian Workshop on Semantic Web Applications and Perspectives (SWAP 2007), Bari, Italy, 18-20 December 2007.

[Col08] S. Colantonio, M. Martinelli, D. Moroni, O. Salvetti, "Decision Support and Image & Signal Analysis in Heart Failure - A Comprehensive Use Case", International Conference on Health Informatics (HEALTHINF 2008), Funchal, Madeira, Portugal, 28-31 January 2008.

[Erika] The Erika Enterprise real-time kernel: http://erika.tuxfamily.org

[Heart] "HEARTFAID: A knowledge based platform of services for supporting medical-clinical management of heart failure within elderly population", European project IST-2005-027107, URL: http://lis.irb.hr/heartfaid/.

[Mal04] D. Malan et al., "Codeblue: An Ad Hoc Sensor Network Infrastructure for Emergency Medical Care", Proc. of the International Workshop on Wearable and Implantable Body Sensor Networks, 2004.

[Mar01] T. Martin and D. Siewiorek. "Non-ideal battery and main memory effects on CPU speed-setting for low power". IEEE Transactions on VLSI Systems, 9(1):29–34, 2001.

[Mar11] M. Marinoni, M. Bambagini, F. Prosperi, F. Esposito, G. Franchino, L. Santinelli and G. Buttazzo, "Platform-aware Bandwidth-oriented Energy Management Algorithm for Real-Time Embedded Systems", Proc. of the 16th IEEE on Emerging Technologies and Factory Automation (ETFA 2011), Toulouse, France, September 2011..

[Mobi] "MOBIHEALTH: Mobile Health Care", European project, IST-2001-36006, URL: http://www.mobihealth.org/.

[Ren10] Y. Ren, R. W. N. Pazzi, A. Boukerche, "Monitoring patients via a secure and mobile healthcare system", IEEE Wireless Communications, Vol. 17, No. 1, February 2010.

[San10] T. Santra, "Mobile Health Care System for Patient Monitoring", in Communications in Computer and Information Science, Springer, Vol. 101, pp. 695-700, 2010

[Sha90] L. Sha, R. Rajkumar, and J. P. Lehoczky, "Priority Inheritance Protocols: An Approach to Real-Time Synchronization", IEEE Transactions on Computers, Vol. 39, No. 9, September 1990.

[Sha09] R. Shahriyar, M. F. Bari, G. Kundu3, S. I. Ahamed, and M. M. Akbar, "Intelligent Mobile Health Monitoring System (IMHMS)", International Journal of Control and Automation, Vol.2, No.3, September 2009.

[Sta05] J. A. Stankovic, I. Lee, A. Mok, R. Rajkumar, "Opportunities and Obligations for Physical Computing Systems," IEEE Computer, 38(11), November 2005.

[Tra09] C. V. Trappey, A. J. C. Trappey, C. S. Liu, "Develop Patient Monitoring and Support System Using Mobile Communication and Intelligent Reasoning", Proceedings of the 2009 IEEE International Conference on Systems, Man, and Cybernetics, San Antonio, TX, USA - October 2009.

[Vir02] G. Virone, N. Noury and J. Demongeot, "A system for automatic measurement of circadian activity deviation in telemedicine," IEEE Transactions on Biomedical Engineering, vol. 49, no. 12, pp. 1463-1469, Dec. 2002.

[Vir06] G. Virone, A. Wood, L. Selavo, Q. Cao, L. Fang, T. Doan, Z. He, and J.A. Stankovic, "An Assisted Living Oriented Information System Based on a Residential Wireless Sensor Network," Proc. of the Transdisciplinary Conference on Distributed Diagnosis and Home Healthcare (D2H2), Arlington, VA, April 2-4, 2006.

[Vir07] G. Virone, T. Doan, A. Wood, and J. A. Stankovic, "Dynamic Privacy in Assisted Living and Home Health Care," in *Joint Workshop On High Confidence Medical Devices, Software, and Systems (HCMDSS) and Medical Device Plug-and-Play (MD PnP) Interoperability*, 2007.

[Woo08] A. Wood, J. A. Stankovic, G. Virone, L. Selavo, Z. He, Q. Cao, T. Doan, Y. Wu, L. Fang, R. Stoleru, "Context-Aware Wireless Sensor Networks for Assisted-Living and Residential Monitoring," IEEE Network, 22(4):26–33, July-August 2008.

[Wu07] K. Wu, X. Wu, "A Wireless Mobile Monitoring System for Home Healthcare and Community Medical Services", The 1st International Conference on Bioinformatics and Biomedical Engineering (ICBBE 2007), Wuhan, China, 6-8 July 2007.