MPSoC ECG Biochip: A Multiprocessor System-on-Chip for Real-Time Human Heart Monitoring and Analysis

Iyad Al Khatib
IMIT, ICT, KTH
Royal Institute of Technology
Stockholm, Sweden
+4687904111
iyad@imit.kth.se

Davide Bertozzi
ENDIF
University of Ferrara
Ferrara, Italy
+390532974832
dbertozzi@ing.unife.it

Francesco Poletti, Luca Benini
DEIS
University of Bologna,
Bologna, Italy
+390512093782
{fpoletti, lbenini}@deis.unibo.it

Axel Jantsch
IMIT, ICT
Royal Institute of Technology (KTH)
Stockholm, Sweden
+4687904124
axel@imit.kth.se

Mohamed Bechara
Hasan Khalifeh, Mazen Hajjar
ECE, FEA
American University of Beirut Beirut, Lebanon
+9613051285
{mfb05, hik05, mah39}@aub.edu.lb

Rustam Nabiev, and
Sven Jonsson
Biomedical Engineering Dept.
Karolinska University Hospital , Sweden
+46858586288
rustam.nabiev@karolinska.se

ABSTRACT
The interest in high performance chip architectures for biomedical applications is on the rise. Heart diseases remain by far the main cause of death and a challenging problem for biomedical engineers to monitor and analyze. Electrocardiography (ECG) is an essential practice in heart medicine, which faces computational challenges, especially when 12 lead signals are to be analyzed in parallel, in real time, and under increasing sampling frequencies. Another challenge is the analysis of huge amounts of data that may grow to days of recordings. Nowadays, doctors use eyeball monitoring of the 12-lead ECG paper readout, which may seriously impair analysis accuracy. Our solution leverages the advance in multi-processor system-on-chip architectures, and is centered on the parallelization of the ECG computation kernel. It improves upon state-of-the-art mostly for its capability to perform real-time analysis of input data, leveraging the computation horsepower provided by many concurrent DSPs, more accurate diagnosis of cardiac diseases, and prompter reaction to abnormal heart alterations. The design methodology to go from the 12-lead ECG application specification to the final hardware/software architecture, modeling, and simulation is the focus of this paper. Our system model is based on industrial components. The architectural template we employ is scalable and flexible.

Categories and Subject Descriptors
C.3 [**Special-Purpose and Application-Based Systems**]: Microprocessor/microcomputer applications, Real-time and embedded systems

CF'06, May 3-5, 2006, Ischia, Italy.

General Terms
Performance, Design, Experimentation.

Keywords
Multiprocessor System-on-Chip, electrocardiogram algorithms, real-time analysis, hardware space exploration

1. INTRODUCTION
Despite the ongoing advances in heart treatment, in the United States [1] and Canada [2] as well as in many other countries, the various forms of cardiovascular disease (CVD) and stroke remain by far the number one cause of death for both men and women of all ethnic backgrounds. According to the World Health Report in 2003, 29.2% of total global deaths are due to CVD, many of which are preventable by action on the major primary risk factors and with proper monitoring [1]. It is estimated that by 2010, CVD will be the leading cause of death in developing countries. Since the rate of hospitalization increases with age for all cardiac diseases [3], a periodic cardiac examination is recommended. Hence, more efficient methods of cardiac diagnosis are desired to meet the great demand on heart examinations. However, state-of-the-art biomedical equipment for heartbeat sensing and monitoring lacks the ability to provide large-scale analysis and remote, real-time computation at the patient's location. The intention of this work is to use multi-processor System-on-Chip (MPSoC) microelectronic solutions to meet the growing demand for telemedicine services, especially in the mobile environment. The project attempts to address the existing problem of reducing the costs for hospitals/medical-centers through using MPSoC designs that may replace biomedical machines and have higher quality, reduce the nurse's and doctor's work-load, and improve quality of care for patients suffering from heart diseases by exploring one potential solution. The proposed solution resolves fundamental mobility problems of patients in a unique way. It

also addresses usability, security and safety of the patients in emergency situations and long-term treatments. From the hospital side, deploying this solution will further reduce the costs of rehabilitating and following up of patients "primary care versus home care", and can result in enhancing effectiveness and proactive planning and decision making by healthcare staff. Home-care ensures continuity of care, reduces hospitalization costs, and enables patients to have a quicker return to their normal life styles. From a technical viewpoint, real-time processing of ECG data would allow a finer-granularity analysis with respect to the traditional eyeball monitoring of the paper ECG readout. Eventually, warning or alarm signals could be generated by the monitoring device and transmitted to the healthcare center via telemedicine links, thus allowing for a prompter reaction of the medical staff. In contrast, heartbeat monitoring and data processing are traditionally performed at the hospital, and for long monitoring periods a huge amount of collected data must be processed offline by networks of parallel computers. New models of healthcare delivery [2] are therefore required, improving productivity and access to care, controlling costs, and improving clinical outcomes. This poses new technical challenges to the design of biomedical ECG equipment, calling for the development of new integrated circuits featuring increased energy efficiency while providing higher computation capabilities.

The fast evolution of biomedical sensors and the trend in embedded computing are progressively making this new scenario technically feasible. Sensors today exhibit smaller size, increased energy efficiency and therefore prolonged lifetimes (up to 24 hours) [4], higher sampling frequencies (up to 10 kHz for ECG) and often provide for wireless connectivity. Unfortunately, a mismatch exists between advances in sensor technology and the capabilities of state-of-the-art heart analyzers [5][6][7]. They cannot usually keep up with the data acquisition rate, and are usually wall-plugged, thus preventing for mobile monitoring. On the contrary, the deployment of wearable devices such as Systems-on-Chip has to cope with the tight power budgets of such devices, potentially cutting down on the maximum achievable monitoring period. In this paper we propose a wearable multi-processor biomedical-chip (MPSoC ECG Biochip) paving the way for portable real-time electrocardiography applications targeting heart disorders. The Biochip leverages the computation horsepower provided by many (up-to-twelve) concurrent DSPs and is able to operate in real-time while performing the finest granularity analysis as specified by the ECG application. Moreover, in case of heart failure emergency aid should arrive in a period of few minutes from the time when the heart failed, otherwise brain damage may occur. Hence, real time analysis must be done in few seconds to allow the alarm signal to reach the emergency aid team, which should act immediately.

The Biochip system builds upon some of the most advanced industrial components for MPSoC design (multi-issue VLIW DSPs, system interconnect from STMicroelectronics, and commercial off-the-shelf biomedical sensors), which have been composed in a scalable and flexible platform. Therefore, we have ensured its reusability for future generations of ECG analysis algorithms and its suitability for porting of other biomedical applications, in particular those collecting input data from wired/wireless sensor networks [8].

The paper goes through all the steps of the design methodology, from application functional specification to hardware definition and modeling. System performance has been validated through functional, timing accurate simulation on a virtual platform. A 0.13µm technology homogeneous power estimation framework leveraging industrial power models is used for power management considerations [9][10].

We point out the need for simulation abstractions matching the application domain, for memory allocation. In addition, with a solution like the MPSoC ECG Biochip, there is potential impact that mobile real-time processing has on the traditional paradigms of healthcare delivery [2][3].

2. MEDICAL BACKGROUND

The electrocardiogram (ECG) is an electrical recording of the heart activity that is used as a diagnosis tool by physicians and doctors to check the status of the heart. The most commonly used way to detect the heart status is the 12-lead ECG technique. This technique uses nine sensors on the patient's body (Fig. 1). The three main sensors are distributed by placing one sensor on the left arm (LA), a second sensor on the right arm (RA), and a third sensor on the left leg (LL). The right leg (RL) is connected by only a wire to be used as ground for the interconnected sensors. By only having these three sensors physicians can use a method known as the 3-lead ECG, which suffers from the lack of information about some parts of the heart but is useful for some emergency cases to have quick analysis. In this respect, medical doctors require more sensors (i.e. more leads). Hence, six more sensors (V1-V6) are added on the chest (Fig. 1). The voltages V1-V6 are measured with respect to ground (G) on the right leg (RL). In some cases, physicians use these sensors to analyze the heart with the 6-lead ECG.

By using all the nine sensors and interconnecting them for the 12-lead ECG gives twelve signals known in biomedical terms as: Lead I, Lead II, Lead III, aVR, aVL, aVF, V1, V2, V3, V4, V5, and V6 (Fig.1-a). The 12-lead ECG produces huge amounts of data especially when used for a long number of hours. Physicians use the 12-lead ECG method, because it allows them to view the heart in its three dimensional form; thus, enabling detection of any abnormality that may not be apparent in the 3-lead or 6-lead ECG technique. Figure 1-b shows an explanatory example of a typical ECG signal. The most important points on the ECG signal are the labeled peaks: P, Q, R, S, T, and U. Each of these peaks is related to a heart action that is of importance for analysis. Figure 1-c shows real recorded signals from 12-leads, which are printed on the eyeballing paper. This paper printout is the classical medical technique used for looking at ECG signals, and it is still used. However, the eyeballing paper print makes the check of the different heart peaks and rhythms difficult and inaccurate due to its dependence on the physician's eyes.

On the other hand, when using digital recording and filtering we can determine the peaks more accurately. Consequently, we can use digital computing to process the sensed data and analyze the heart beat. Figure 4-a shows the ECG signal of a normal heart that was recorded digitally for 5 seconds with a sampling frequency of 250Hz. The labeled peaks (P, Q, R, S, T, U), and the time intervals between them can show if the heartbeat is healthy or unhealthy. In addition, there are normal medical ranges for the inter-peak time intervals, and every combination of different

inter-peak intervals proves a type of heart illness. The most important of the peaks is the R peak, which refers to the largest heart blood pump.

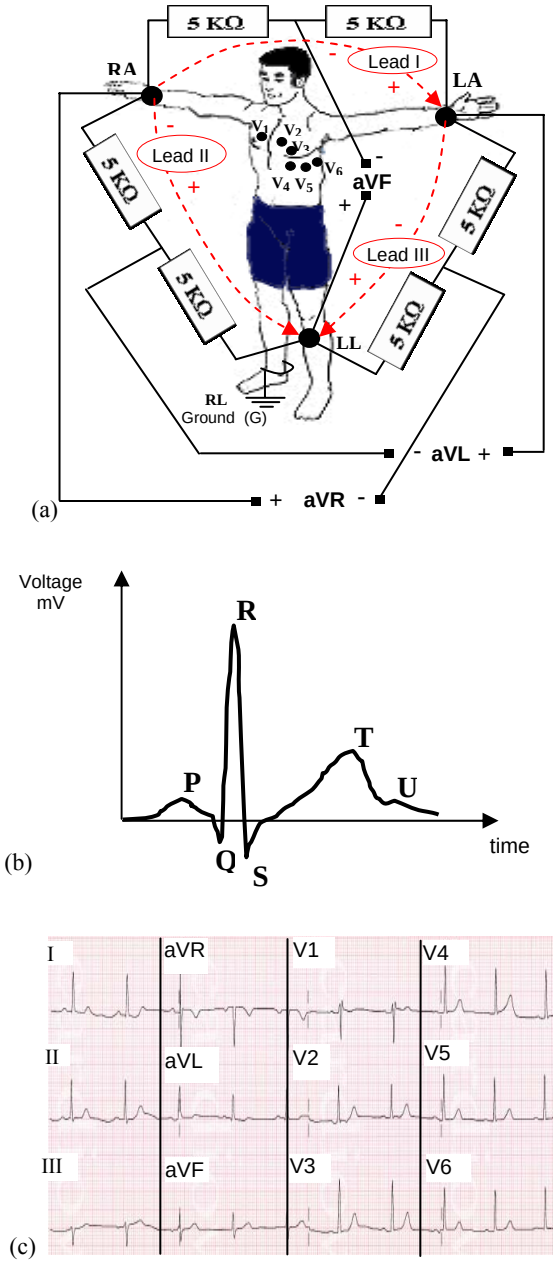

Figure 1. 12-lead ECG: (a) Sensors on a human body, RA is right arm sensor, LA is left arm sensor; LL is left leg, RL is right leg that is grounded (G); (b) Example of a typical healthy ECG signal for leads I and II. (c) Complete paper readout, which is not accurate to see peaks nor easy to read for long recordings. The 12 lead signals are: Lead I, Lead II, Lead III, aVR, aVL, aVF, V1-G, V2-G, V3-G, V4-G, V5-G, and V6-G.

3. PREVIOUS WORK

ECG monitoring and analysis have been commercially explored in many companies and research organizations. However, we are not aware of any single-chip real-time analysis solution for full 12-lead ECG, which is able to accurately study the heart rhythmic period and can diagnose all the peaks: P, Q, R, S, T and U and their inter-peak intervals to result in a disease diagnosis.

Most of the work done involves only recording of huge amounts of data in large storage media and then analyses, but not allowing the ease of patient mobility. Most of the time, the patient has to be confined to a bed for a number of hours. Some commercial solutions are only capable of concluding if the heart beat is normal or abnormal but can not specify the period nor the disease. Other real time solutions available in the market, in healthcare institutes, and in research organizations, are only able to sense and transmit ECG data [11] to: either a local machine [12] or to a distant healthcare center [13]. In both cases, real-time analysis is not available. Moreover, the commercial solutions under study [14] do not look into the parallelization of the ECG analysis into multiple cores, so to speed up processing.

4. SYSTEMS AND ALGORITHM

ECG analysis requires three main phases: (i) signal-reading from the leads, (ii) filtering the lead-signal, and finally (iii) analysis (Fig.2). Firstly, the signal sensing phase requires an A/D converter in order to be able to have digital data for our digital filter. We use 16 bit A/D converters, because our analysis algorithm and ECG Biochip are designed based on having 16-bit filtered data as input. We briefly discuss the filtering method we use as an essential part of our proposed solution, and then we discuss the ECG Biochip design that depends on this filtering step.

4.1 Input Data Filtering

In general, data provided by biomedical sensors suffers from several types of noise that are due to the specific sensors used (DC-offset is one example), patient movements, and environment interference like other frequencies in the air [15].

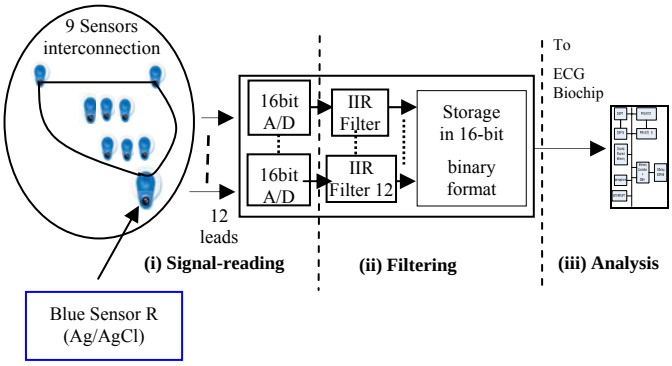

Figure 2. The System for sensing and filtering of ECG lead signals before sending data to the ECG Biochip for analysis. Blue Sensor R is form Ambu Inc. [4]. The interconnection of the 9 sensors is shown in Fig. 1.

We use state of the art commercial sensors from Ambu Inc. silver/silver chloride "Blue Sensor R" [4] shown in Fig. 2. It is characterized by: 24 hour lifetime, superior adhesion, optimal signal measuring during stress tests. It is small to carry (57mm x 48mm).

This reality of having many noise factors may lead to inaccurate ECG analysis, for example to detect 2 peaks when there should be only one. See for instance the R-Peak detection marked by circled areas in Fig.3. To overcome all the problems related to sensor noise affecting our analysis, we designed an IIR filter with order 3 that outputs its results in 16-bit binary format (Fig.2). Another main advantage of using the IIR filter is to eliminate the noise that is directly proportional to the DC offset of the sensed ECG [15], which is around 0.1mV. The two plots in Fig.3 clearly show how the filtering algorithm remedies this problem. In our implementation, the filter is implemented in hardware on a dedicated chip feeding the external SDRAM memory of our ECG Biochip.

Our filter is the convolution of the noisy signal with the filter impulse response given by (1):

$$y[n] = \sum_{k} h[k] \times x[n-k] \qquad (1)$$

where, $x[n]$ is the noisy signal, $h[n]$ is the filter impulse response, and n is the samples index. This filter in (1) is also an infinite impulse response (IIR, *Chebychev filter*), so it can be written as (2):

$$y[n] = \sum_{l=0} x[n-l] \times b[l] - \sum_{m=1} y[n-m] \times a[m] \qquad (2)$$

where, y is the output of the filter and x is the input, b is the vector that contains the filter coefficients for signal x, and a is the vector that contains the filter coefficients for output y. The upper limits of the coefficients are dependent on the order of the filter being used. Our IIR filter is of order 3, because our ECG data does not require higher orders. We can improve our filter (when needed) by simply knowing the needed values of the coefficients in vectors $a[.]$ and $b[.]$.

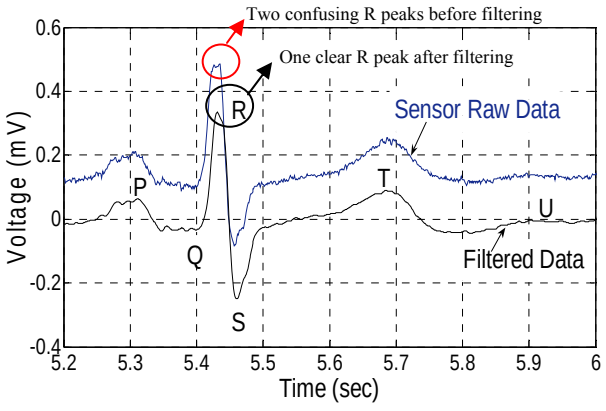

Figure 3. ECG raw and filtered data.

4.2 The Algorithm

The proposed algorithm was conceived to be parallel and hence scalable from the ground up. Since each lead senses and analyzes data independently, each lead can then be assigned to a different processor. So, to extend ECG analysis to 15-lead ECG for example or more, then what is required is to just change the number of processing elements in the system. The program reads a data file in chunks of four seconds. We discuss below the reason for the choice of the 4-second chunks. The data file mainly holds the values of the ECG at the lead in binary format. So by reading the data continuously every 4 seconds, we would be emulating a real sensor sending continuous data to an intermediate buffer that holds 4 seconds of data sampled at a certain frequency, typically 1000 Hz. We used an autocorrelation function based-methodology to calculate the period of the heartbeat since it gives more accurate results than the conventional methods searching for the distance between two peaks. These latter methods are only effective to get the period for a normal person, while our technique is able to detect it even in case of abnormalities. We validated our algorithm over several different input traces [16] and medical scenarios.

$$R_y[k] = \sum_{n=-\infty}^{n=\infty} y[n] \times y[n-k] \qquad (3)$$

where, Ry is the autocorrelation function, y is the signal under study, n is the index of the signal y, and k is the number of lags of the autocorrelation. We run the experiments for n = 1250, 5000 and 50,000 relative to the sampling frequencies of 250, 1000, and 10,000Hz, respectively. Figure 4-c shows the period of the heartbeat as the distance between two consecutive peaks of the autocorrelation function of the derivative (Fig.4-b) of the input signal (Fig.4-a). In order to be able to analyze ECG data in real-time and to be reactive in transmitting alarm signals to healthcare centers (in less than 1 minute), a minimum amount of acquired data has to be processed at a time without losing the validity of the results. For the heart beat period, we need at least 4 seconds of ECG data in order for the ACF to give correct results.

In addition, for both cases of healthy and unhealthy hearts, the 4-seconds granularity contains at least 5 heart cycles that are needed for ECG analysis. The algorithm differs from one lead (DSP) to another in the calculation method (code) of the peaks and inter-peak intervals done in process 2. Although the autocorrelation function in (3) is the same for all leads and thus all DSPs, each DSP still has to process the autocorrelation of the derivative of its input signal (relative ECG-lead filtered-data). Therefore, it is essential to optimize our time period for this algorithm, because it affects the very high number of multiplications (around 1.75 million multiplications) that- in turn- affect the Biochip performance. For this reason, we do not go beyond 4 seconds, which is a critical number for our Biochip performance, hence its application-specific design.

The autocorrelation function is deployed within the algorithm shown in Fig.5, which computes the required medical parameters: heart period, peaks P, Q, R, S, T, & U, and inter-peak time spans. Peak heights and inter-peak time ranges outside normal values, which indicate different kinds of diseases, are detected with our algorithm.

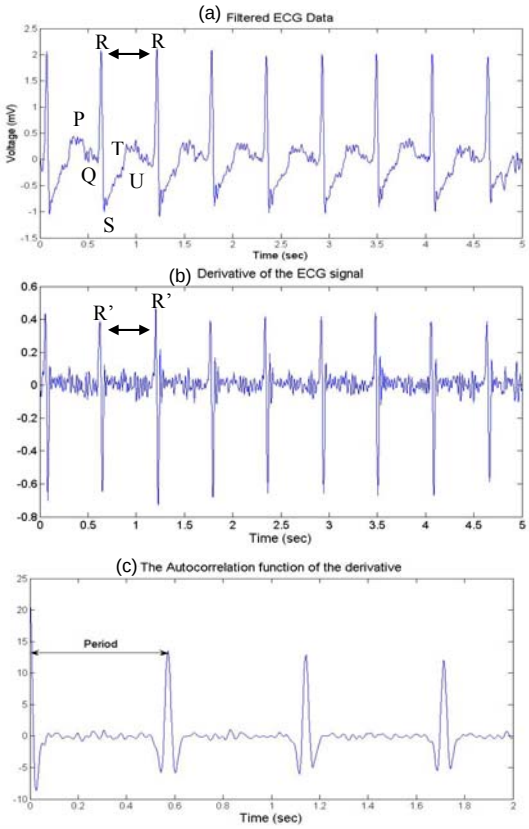

Figure 4. Heart period analysis for ECG Lead I, which is processed by DSP 1 on the Biochip: (a) signal with peaks P, Q, R, S, T, and U; (b) derivative amplifying the R peaks that we label as R' peaks; (c) autocorrelation of the derivative with clear significant periodic peaks.

Each DSP on the biochip will have to run this algorithm. Each lead has its specific ECG plot and characteristic as shown in Fig.1-c. From a functional viewpoint, the algorithm consists of two separate execution flows: one that finds the period using the autocorrelation function (process 1 in Fig.5), and another that finds the number, amplitude and time interval of the peaks in the given 4-second ECG data (process 2 in Fig.5). In process 1, we firstly find the discrete derivative of the ECG signal. This will not affect the analysis since the derivative of a periodic signal is periodic with the same period. The advantage of taking the derivative, and thus adding some overhead to the code, is that the fluctuations taking place in the signal and especially those around the peaks would converge to zero. Moreover, the time consumption of the code part related to calculating the derivative of the ECG signal is negligible compared to the rest of the algorithm code, especially the autocorrelation part (containing millions of multiplications). After the derivative of the ECG filtered data, we run the autocorrelation of the derivative signal which is, by definition, periodic with the same period as that of the ECG signal under study. The autocorrelation function will start at a maximum point, go down to zero, and then rise to another peak. The index of this peak represents the period of the signal. While we are performing the derivative and the autocorrelation, the values of these functions are dumped in 16-bit binary format. In process 2, a threshold is used to find the peaks. This threshold is 60% of the highest peak in the given interval. This was the choice after performing several experiments with different real ECG data.

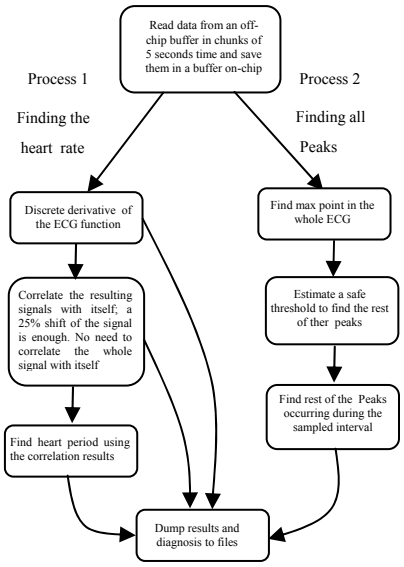

Figure 5. Algorithm for analyzing 12-lead ECG data for 250Hz, 1000Hz, and 10,000Hz sampling frequencies.

5. ARCHITECTURE DEFINTION

In order to process filtered ECG data in real-time, we chose to deploy a parallel Multi-Processor System-on-Chip architecture. The key point of these systems is to break up functions into parallel operations, thus speeding up execution and allowing individual cores to run at a lower frequency with respect to traditional monolithic processor cores.

Technology today allows the integration of tens of cores onto the same silicon die, and we therefore designed a parallel system with up to 13 masters and 16 slaves (see Fig.6). Since we are targeting a platform of practical interest, we chose advanced industrial components [17]. The processing elements are multi-issue VLIW DSP cores from STMicroelectronics, featuring 32kB instruction and data caches.

The speeds of the processors are set to 200MHz, which is the speed of the bus and the memory. These cores leverage the flexibility of programmable cores and the computation efficiency of DSP cores. These features allow the reuse of this platform for other biomedical applications other than the 12-lead ECG, thus making it cost-effective. Each processor core has its own private memory (512kB each), which is accessible through the bus, and can access an on-chip shared memory (8kB are enough for this application) for storing computation results. Other relevant slave components are a semaphore slave, implementing the test-and-set operation in hardware and used for synchronization purposes by the processors or for accessing critical sections, and an interrupt

slave, which distributes interrupt signals to the processors. Interrupts to a certain processor are generated by writing to a specific location mapped to this slave core. The STBus interconnect from STMicroelectronics was instantiated as the system communication backbone. STBus can be instantiated both: as a shared bus or as a partial or full crossbar, thus allowing efficient interconnects design and providing flexible support for design space exploration. In our first implementation, we target a shared bus to reduce system complexity (see Fig.6) and assess whether application requirements can already be met or not with this configuration. We then explore also a crossbar-based system, which is sketched in Fig.7. The inherent increased parallelism exposed by a crossbar topology allows decreasing the contention on shared communication resources, thus reducing overall execution time. In our implementation, only the instantiation of a 3x6 crossbar was interesting for the experiments. We put a private memory on each branch of the crossbar, which can be accessed by the associated processor core or by a DMA engine for off-chip to on-chip data transfers. Finally, we have a critical component for system performance which is the memory controller. It allows efficient access to the external 64MB SDRAM off-chip memory. A DMA engine is embedded in the memory controller tile, featuring multiple programming channels. The controller tile has two ports on the system interconnect: one slave port for control and one master port for data transfers. The overall controller is optimized to perform long DMA-driven data transfers. Embedding the DMA engine in the controller has the additional benefit of minimizing overall bus traffic with respect to traditional standalone solutions. Our implementation is particularly suitable for I/O intensive applications such as the one we are targeting in this work.

In the above description, we have reported the worst case system configurations. In fact, fewer cores can be easily instantiated if needed. In contrast, this architectural template is very scalable and allows for further future increase in the number of processors. This will allow to run in real time even more accurate ECG analyses for the highest sampling frequency available in sensors (10,000Hz, and 15 leads, for instance), since this platform is able to provide scalable computational power.

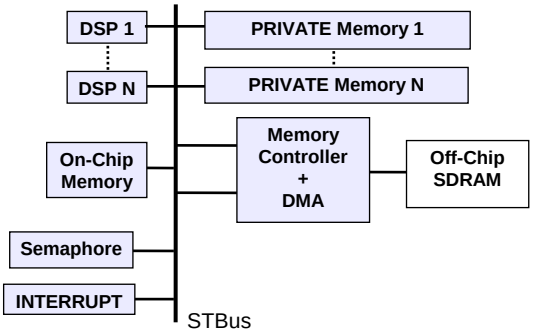

Figure 6. Single bus architecture with STBus interconnect. In our experiments max. N = 12. The solution is in general scalable since we can increase the number of leads to analyze by just increasing N.

The entire system has been simulated by means of the MPSIM simulation environment [17], which provides for cycle-accurate functional simulation of complete MPSoCs. The simulator provides also a power characterization framework leveraging 0.13µm technology-homogeneous industrial power models from STMicroelectronics [9][10].

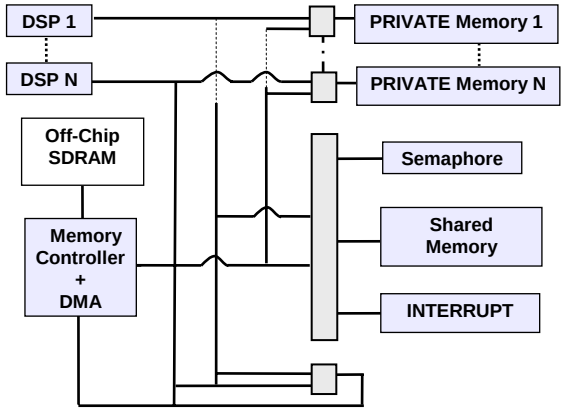

Figure 7. Full crossbar architecture with STBus interconnect. In our experiments max. N = 12. The solution is in general scalable since we can increase the number of leads to analyze by just increasing N.

We believe that for life-critical applications, low-level accurate simulation is worth doing, in order to perfectly understand system level behaviour and have a predictable system with minimum degrees of uncertainty.

Each processor core programs the DMA engine to periodically transfer input data chunks onto their private on-chip memories. Moved data corresponds to 4 seconds of data acquisition at the sensors: 10kB at 1000Hz sampling frequency, transferred on average in 319279 clock cycles (DMA programming plus actual data transfer) on a shared bus with 12 processors. The consumed bus bandwidth is about 6MBytes/sec, which is negligible for an STBus interconnect, whose maximum theoretical bandwidth with 1 wait state memories exceeds 400Mbyte/sec. Then each processor performs computation independently, and accesses its own private memory for cache line refills. Different solutions can be explored, such as processing more leads onto the same processor, thus impacting the final execution time. Output data, amounting to 64 bytes, are written to the on-chip shared memory, but their contribution to the consumed bus bandwidth is negligible. In principle, when the shared memory is filled beyond a certain level, its content can be swapped by the DMA engine to the off-chip SDRAM, where the history of 8 hours of computation can be stored. Data can also be remotely transmitted via a telemedicine link.

6. EXPERIMENTS AND RESULTS

We ran several different comparisons to test the functionality and performance of our algorithm. The first analysis was done to profile the execution of the code and to determine the best coding solution in terms of energy, execution time, and precision. Furthermore, we have explored the design space searching for the

best platform configuration for the 12-lead ECG data analysis. Alternative system configurations have been devised for different levels of residual battery lifetime, trading off power with accuracy.

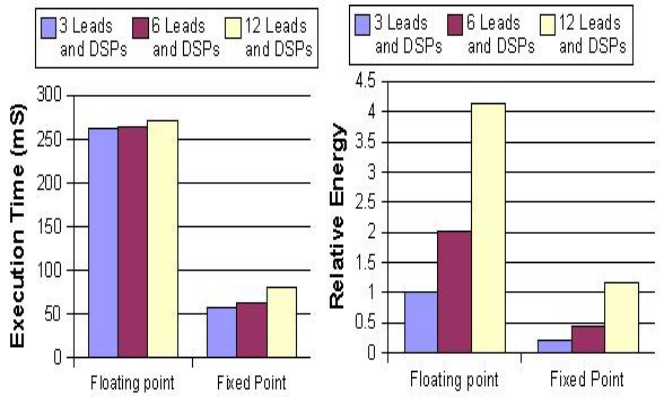

Figure 8. Comparison between different code implementations for analysis of the 3-lead, 6-lead and 12-lead ECG. Data analysis for each lead is computed in a separate core. Sampling frequency of input data was 250 Hz.

Figure 8 shows the results for two different code implementations: the first one relies on floating point data types, while the second one uses fixed point data [18] with an exponent of 22. We have performed the analysis for 3, 6 and 12 leads; furthermore we process each lead on a separate core. We found that the precision of the results obtained with fixed point code, by using 64 bit integer data types representation, almost matches the results obtained with floating point code for a large number of input data traces. On the contrary, the time needed to process data, and also the energy required, decreases up to 5 times. This is mainly due to the fact that, like many commercial DSPs, our processor cores do not have a dedicated floating point unit. Therefore floating point computations are emulated by means of a C software library linked at compile time. Fig.8 also shows that even with 12 concurrent processors, the bus is not saturated, since we observe negligible effects on the stretching of task execution times. In contrast, adding more processors determines a linear increase in energy dissipation.

In the second experiment, we were looking for the best platform configuration to meet the time constraints of the application and reduce energy consumption. We considered the scenario with 12 leads ECG data and 1000Hz sensor sampling frequency (the nowadays medical reference frequency for ECG). We have scaled the number of cores from 12 to 2 processors, in an attempt to minimize system resources while meeting application requirements. In the 12-processor configuration we computed each lead on a separated DSP. As a result, the total execution time to perform a 4 second period analysis took 1.1 seconds as shown in Fig. 9.

This means that we are able to meet the 4-second deadline, after that a new set of input data has to be processed. The slack in this case is so large that we opted for processing 2 leads on each DSP, observing an execution time degradation of 2 times (computation in 2.2sec).

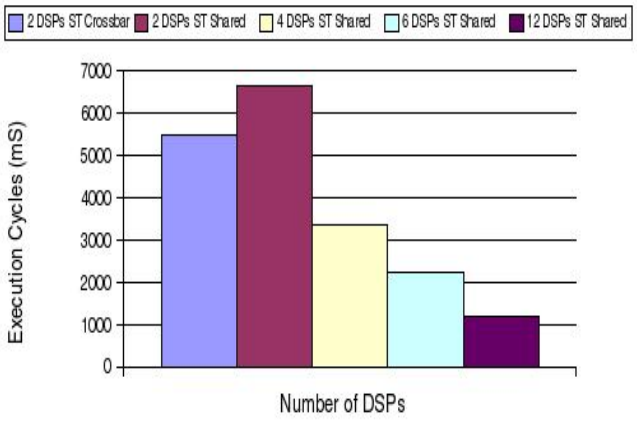

Figure 9. Time analysis for 12-lead ECG data with 1000Hz sensor sampling frequency using different numbers of DSP cores: 2, 4, 6, and 12 processing cores, with different interconnects.

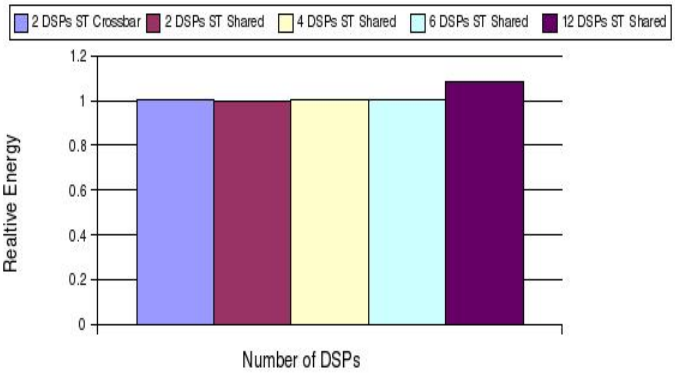

Figure 10. The relative energy consumption for 1000Hz sensor sampling frequency for ECG data using different numbers of DSP cores: 2, 4, 6, and 8 processing cores, with different interconnects.

Pushing our design to check for minimizing the number of DSPs and keeping the full algorithm (for its accuracy and ability of complete analysis), the solution with 4 DSPs (i.e. 3 lead algorithms on each processor with 1000Hz sampling frequency) turns out to be the first design able to complete the computation reasonably earlier than 4sec. The remaining time is very useful in critical cases, since it can be used to trigger an emergency procedure; results for the computations of the different leads can be correlated and- eventually- an alarm is transmitted to a remote healthcare center via telemedicine links. Finally, we explore the solution with two DSPs and two different communication architectures: *a shared bus* (Fig.6), and *a full crossbar system* with STBus (see Fig. 7). Even though the full crossbar allows us to decrease the execution time by almost 10%, for both architectures the execution time is bigger then the 4sec deadline, and do not therefore represent viable options.

Interestingly, Fig.10 shows energy dissipation for all considered system configurations. The energy is almost unchanged, since employing a lower number of processor cores decreases the number of power-contributing cores, but increases execution time. The two effects counterbalance each other.

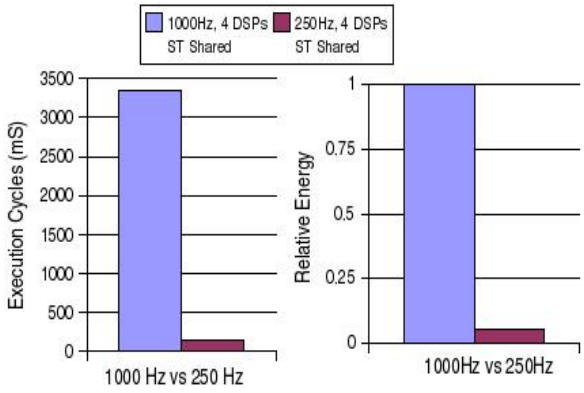

Figure 11. Results of the 12-lead analysis with 4 DSPs.

Hence, we decided to adopt the 4-DSP core solution for the 12 lead ECG analysis. We now explore how the energy consumption scales by increasing the sampling frequency of the sensor. A higher sampling frequency increases the precision of the results. However, unfortunately, the complexity of the algorithm scales exponentially. In fact, Fig. 11 shows that the time taken for the ECG analysis in the 250Hz case is 21 times lower than in the 1000Hz case. Moreover, the energy consumption in the 250Hz case is 90% lower than in the 1000Hz case. So, in cases of emergency, and if the battery is low, a down sampling of the input data from 1000Hz to 250Hz will be good to assure a longer lifetime for the battery to keep the 12 lead analysis running, hence paying on the frequency to win time and power.

7. CONCLUSION

In this paper, we present a novel ECG Biochip solution leveraging the computation horsepower of many concurrent DSP cores to process ECG data in real-time. This solution paves the way for novel healthcare delivery scenarios (e.g., mobility) and for accurate diagnosis of heart-related diseases. We described the design methodology for the MPSoC and explored the configuration space looking for the most effective solution, performance- and energy-wise.

8. REFERENCES

[1] Fuster, V., Epidemic of Cardiovascular Disease and Stroke: The Three Main Challenges, *Circulation, Vol. 99, Issue 9*, March 1999, pp. 1132-1137.

[2] Heart and Stroke Foundation of Canada, *The Changing Face of Heart Disease and Stroke in Canada 2000*, Annal report, 1999.

[3] Chan, C., Han, J., and Ramjeet, D., *LabVIEWTM Design of a Vectorcardiograph and 12-Lead ECG Monitor*, Final Year Project for the Bachelor of Science Degree in the University of Manitoba, March 2003.

[4] Ambu, Inc. *biomedical devices company*, www.ambuusa.com

[5] Harland, C., Clark T., and Prance, R., Electric Potential Probes– New Directions in the remote sensing of the human body, *Measurement Science and Technology*, Vol. 13, 2002, pp. 163-169.

[6] Harland, C., Clark, T., and Prance, R., High resolution ambulatory electrocardiographic monitoring using wrist-mounted electric potential sensors, *Measurement Science and Technology, Vol. 14*, 2003, pp. 923-928.

[7] Malmivuo, J., and Plonsey, R., *Bioelectromagnetism: Principles and Applications of Bioelectric and Biomagnetic Fields*, Oxford University Press, 1995.

[8] Chevrollier, N., and Golmie, N., On the Use of Wireless Network Technologies in Healthcare Environments, In *Proceedings of the fifth IEEE workshop on Applications and Services in Wireless Networks (ASWN 2005)*, June 2005, pp. 147-152.

[9] Loghi, M., Poncino, M., and Benini, L., Cycle-Accurate Power Analysis for Multiprocessor Systems-on-a-Chip, In *GLSVLSI04: Great Lake Symposium on VLSI*, April 2004, pp.401-406.

[10] Bona, A., Zaccaria, V., and Zafalon, R., System level power modeling and simulation of high-end industrial network-on-chip, In *Proceedings of Design and Test in Europe Conference (DATE)*, February 2004, pp. 318-323.

[11] Lo, B., Thiemjarus, S., King R., and Yang, G., Body Sensor Network–A Wireless Sensor Platform for Pervasive Healthcare Monitoring, In *Adjunct Proceedings of the 3rd International Conference on Pervasive Computing (PERVASIVE'05)*, May 2005, pp.77-80.

[12] Association of Cardiac Technology in Victoria, *ACTIV*, http://www.activinc.org.au/

[13] Code Blue, *Wireless Sensor Networks for Medical Care*, http://www.eecs.harvard.edu/~mdw/proj/codeblue/

[14] BIOPAC Systems Inc., http://biopac.com/

[15] Company-Bosch, E., Hartmann, E., ECG Front-End Design is Simplified with MicroConverter, *Journal of Analog Dialogue, Vol. 37*, November 2003.

[16] PhysioBank, *physiologic signal archives, for biomedical research*, http://www.physionet.org/physiobank/database/ptbdb/

[17] Loghi, M., Angiolini, F., Bertozzi, D., Benini L., and Zafalon, R., Analyzing On-Chip Communication in a {MPSoC} Environment, In *Proceedings of the Design and Test in Europe Conference (DATE)*, February 2004, pp.752-757.

[18] ARM DAI 0033A, Note 33, *Fixed Point Arithmetic on the ARM*, September 1996.