Deliverable report for

PHIDIAS
"Ultra-Low-Power Holistic Design for Smart Biosignals Computing Platforms"

Grant Agreement Number 318013

Deliverable D 1.2
Report on structured sparsity for bio- signals

Due date of deliverable: 30/09/2013

Lead beneficiary for this deliverable: EPFL
Contributors:
- **EPFL:** development and performance assessment of joint compressed sampling. Writing of the deliverable
- **ALL:** revision

Dissemination Level:		
PU	Public	X
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	

Version: 1.3 (revision)	Date 04.03.2014
Draft of the WP Leader	
Commented version for amendment	
Version accepted by the Steering Board	X

1. Description of task

This deliverable report the results of the Task 1.2 of WP1, describing the investigated advanced sparsity techniques for smart sensing and optimized recovery of bio-signals. By merging both sampling and compression, CS allows to develop practical ultra-low power read-out systems for wireless bio-signal monitoring devices, where large amounts of sensor data need to be transferred through the power-hungry wireless links.

Building on the work illustrated in D1.1, in this deliverable we propose a novel approach for joint compression of multi-lead ECG signals, where strong correlations exist between leads. The technique can be exploited to optimize the quality of a signal after reconstruction, or, alternatively, minimize the amount of data sent through an energy-hungry wireless link, thus addressing the challenge of ultra-low-power embedded monitoring of multi-lead bio-signals.

Figure 1 exemplifies the reconstruction quality of the proposed technique, plotting an excerpt from original and reconstructed signals. For clarity, only one of the 15 signals composing the recording is presented, and plots are horizontally shifted to separate them. Across the recording, the compression ratio obtained in the example is 90%.

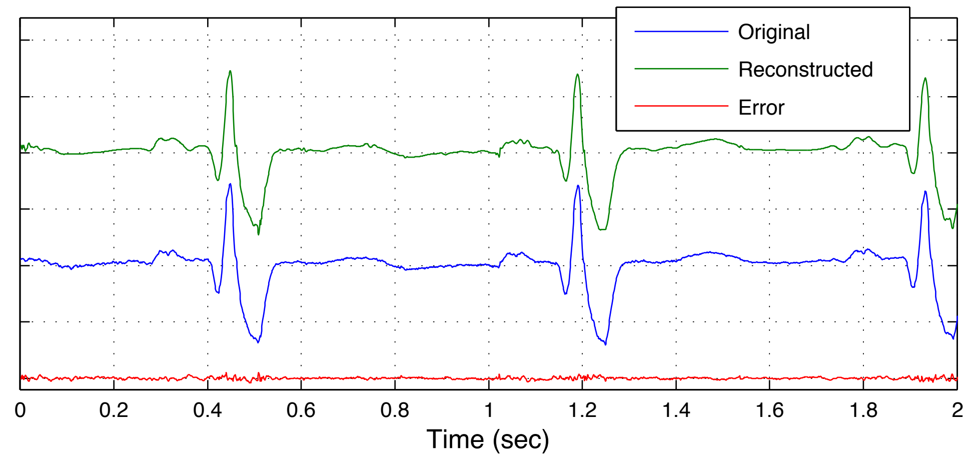

Figure 1: Reconstructed (top) and original (bottom) signal using the joint reconstruction CS technique, with a compression ratio of 90%

2. Joint Compressed sensing of multi-lead ECG

The fundamentals of CS have been discussed in D1.1, which presented experimental evidence of CS performance when inputs are considered in isolation. Nonetheless, in a real scenario many of the bio-signals are acquired on multiple channels. As an example, for the ECG signals we have different standards, e.g.: 3-lead or 12-lead combinations. In these cases, sampling, compressing and reconstructing each signal individually is clearly sub-optimal, because leads are not independent sources, and are in fact strongly correlated. In the case of the ECG acquisitions, all signals can be considered as different projections of a single multidimensional source, which is the electrical field produced by the heart. Thus not only these leads are not independent, but each lead conveying useful information about other leads.

In this deliverable we discuss how mutual information between the leads of bio-signals can be exploited to optimized CS-based compression, considering the example of multi-lead ECG acquisitions. We also propose the techniques to exploit this information for optimized recovery of the joint compressed sensing. This optimization will directly translated to less measurement needed for the recovery, leading to the minimization of transmission bandwidth (and, ultimately, of power consumption) without quality degradation.

3. Compressed sensing and recovery algorithms

3.1. Compressed sensing and background on sparsity

Let's denote the vector x as a real–valued N-dimensional vector of samples of single-lead bio-signal and $\alpha = \Psi x$ is the expansion of the vector x in a given sparse domain Ψ. For example, it is well known that ECG signals have a very sparse representation on the wavelet domain, which means that their wavelet coefficients α are sparse, or more accurately compressible.

By sparse we mean that α has few nonzero elements and rest of the elements are zero (exact sparse), or very close to zero (compressible). The original signal vector x then can be described by linear superposition of S elements of an orthonormal basis in which the signal is sparse: $x \approx \sum_{i=1}^{S} \alpha_i \psi_i$, with $S << N$.

Given the sparsity of the input signal, compressed sensing states that, only $S\log\left(N/S\right)$ linear measurements of the vector x are enough to preserve all information of the signal. This measurement vector is denoted as $y = \Phi\alpha = \Phi\Psi x$. The matrix Φ is called the sensing matrix. In our experiment the sensing matrix is binary and sparse, which is proven to be the time and space optimal choice, and very close to bounds of the measurement optimality, as detailed in D1.1.

Under broad conditions that must be satisfied by x it is then possible to recover the original signal x from its measurements by solving a convex optimization problem. The problem is formulated in the form

$$\min_{\tilde{\alpha} \in \Re^N} \|\tilde{\alpha}\|_1 \text{ subject to } \|\Phi\Psi\tilde{\alpha} - y\|_2 \le \sigma \qquad [1]$$

The recovered signal $\hat{x}$ is then computed as $\hat{x} = \Psi^T \hat{\alpha}$.

The first part of equation [1] is a gradient optimization for minimization of the error and second part adds the sparsity constraint to force to reach the sparse solution. In case of the compressible signals it is proven that the recovered $\hat{x}$ is the best S-term approximation of the original signal x.

It has to be noted that further lossless compression coding can be applied to the signal transformed in the sparse domain. In this deliverable we only report the results of the Compressed Sensing, without compression coding.

As discussed before, in many cases the monitoring ECG signals, and bio-signals in general, is performed by simultaneous recording of multiple leads. Let's indicate $X \in R^{N \times L} = [x_1, x_2, ..., x_L]$ as a matrix of composed of L leads, each row representing a vector of ECG data $x_i \in R^{N \times 1}, i \in \{1, 2, ..., L\}$. The whole matrix X is acquired and compressed at the same time $Y = \Phi X, Y \in R^{M \times L}$.

4. Performance Metrics and Databases

To quantify the compression performance while assessing the diagnostic quality of the compressed ECG records, we employ the two most widely used performance metrics, namely the compression ratio (CR) and percentage root-mean-square difference (PRD). The compression ratio is defined as the percent reduction in number of bits required to represent the signal after compression relative to the original number of bits: $CR = \frac{b_{orig} - b_{comp}}{b_{orig}} \times 100\%$, where b_{orig} and b_{comp} represent the number of bits required for the original and compressed signals, respectively. The percentage root-mean-square difference (PRD), and associated signal-to-noise ratio (SNR), quantifies the percent error between the original signal vector x and the reconstructed $\hat{x}$:

$$PRD = \frac{\|x - \hat{x}\|_2}{\|x\|_2} \times 100$$

$$SNR = -20 \log_{10}(0.01 \times PRD)$$

Table 1 reports the resulting different quality classes and corresponding PRD:

PRD	Reconstructed Signal Quality
0 ~ 2%	"Very good" quality
2 ~ 9%	"Very good" or "good" quality
≥ 9%	Not possible to determine the quality group

Table 1: PRD and corresponding quality class

4.1. Database

We investigated CS of ECG data, obtained from the PTB Diagnostic ECG Database, available on the physionet website[1]. The database contains 549 records from 290 subjects. Each record includes 15 simultaneously measured signals: the conventional 12 leads (i, ii, iii, avr, avl, avf, v1, v2, v3, v4, v5, v6) together with the 3 Frank lead ECGs (vx, vy, vz). Each signal is digitized at 1K samples per second, with 16-bit resolution over a range of ± 16.384 mV.

5. Methods

5.1. Joint Sparsity and support selection

The problem of the CS could be seen as finding the best k-term approximation of the sparse vector α which satisfies the $y = \Phi\Psi\hat{\alpha}$. Like in many greedy algorithms the problem could be seen as a two-folded problem. First, we need to find the support that achieves the best approximation, and then we also need to find the values of the entries in this support that achieves the best approximation.
Let's imagine $\hat{\alpha}$ is the best k-term approximation of the α, then:

$$\hat{\alpha} = \underset{x=\|x\|_0 \le S}{\arg\min} \|x-\alpha\|_2^2$$

Let $B = \sup(\hat{\alpha})$ is the support of the vector that achieves the best approximation.

$$(\hat{\alpha})_B = (\alpha)_B \text{ and } M = (S\log(N/S))$$

Where $(\hat{\alpha})_{B^C}$ is the compliment set of the support B. It can be can proved that:

$$\begin{aligned} B = \sup(\hat{\alpha}) &= \underset{K:|K|\le S}{\arg\min} \|(\alpha)_K - \alpha\|_2^2 \\ &= \underset{K:|K|\le S}{\arg\max}(\|\alpha\|_2^2 - \|(\alpha)_K - \alpha\|_2^2) \\ &= \underset{K:|K|\le S}{\arg\max}(\|(\alpha)_K\|_2^2) \\ &= \operatorname{argmax} \sum_{i\in K} |\alpha_i|^2 \end{aligned}$$

[1] http://www.physionet.org/physiobank/database/ptbdb/

Where K is the set on nonzero elements in the best S-term approximation. So the best support is the set of indices contributing the most to the 2-norm of α. In other words, the best support is achieved by choosing the indices of the largest S elements of α. In case the support set is known, then the CS problem can be solved, if enough measurements are provided.

5.2. Joint support selection

As mentioned before, we propose to leverage the similarities between multiple channels to obtain a high-performance CS-based methodology. Following this intuition, we propose and evaluate a novel algorithm for joint reconstruction of multi-lead bio-signal acquisitions. Figure 2 shows the support of the best S-term approximation of the investigated multi-lead ECG recordings. It shows that support set of the S-term approximation for leads are very similar, so that a joint sparse support selection can be performed in the reconstruction algorithm.

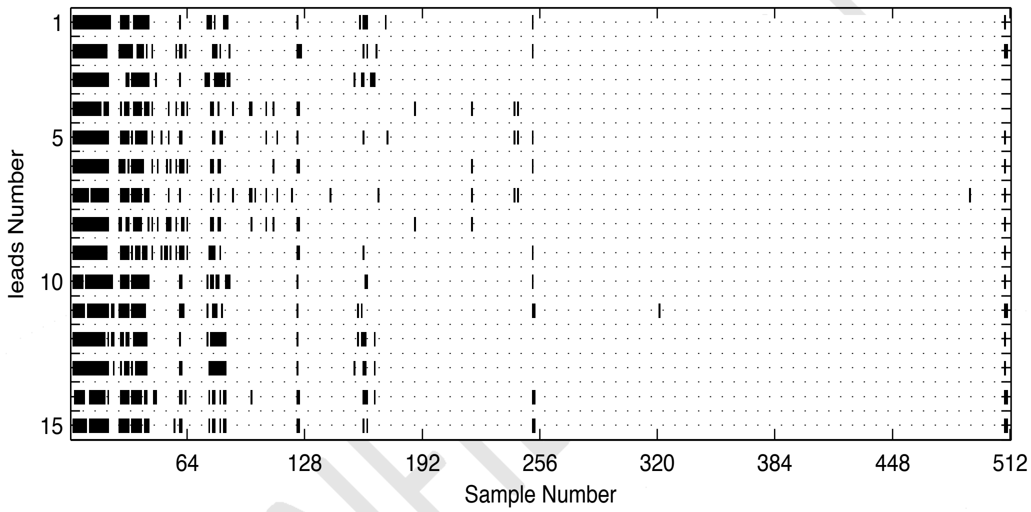

Figure 2: Support of the best S-term approximation of multi-lead ECG signals.

To solve the reconstruction problem we the L1 norm minimization problem detailed in Section 3.1 is done simultaneously for all leads, and S-term approximation of the sparsity support of the first lead is used for recovering all the other leads.

6. Encoding algorithm

The processes of collecting the measurements (compression) of the multi-leads ECG signals follows the routine used for the single lead case. The difference is in the selection of the reference lead. In our study, we always lead 1. This lead is compressed at a lower rate, to guarantee a very good support selection. The extracted support set is then used as an additional constraint in solving the CS problem, where all indices out of the support are set to zero.

7. Results

7.1. Compression ratio and best S-term approximation

D1.1 reported that ECG signals are compressible but they are not exact sparse. This means that their expansion of the proper transformation domain, like the discrete wavelet domain, results in a compressible representation.

In case of exact sparse signals, the number of required measurements for the measurement optimal sensing matrix is $M = O\left(S\log(N/S)\right)$, while for compressible signal the number of measurements depends on the level of sparsity and the desired reconstruction quality, and can be therefore bigger than M .

Moreover, in the illustrated scenario of joint support selection, the amount of similarity between best S-term approximation support of each lead and the reference one plays an important role, and must be investigated. To do so, we have defined a constant value δ as the fraction of sparsity level S over the number of measurements M: $\delta = S/M$. In our simulations for each number of measurements M , different values for δ corresponding to different values of S-term approximations of other leads are investigated. Figure 3 shows the results of output SNR over all records of the database. Different curves for each values of δ are plotted. The plot shows that how for different values of acquired number of the measurements the best S-term approximation will change..

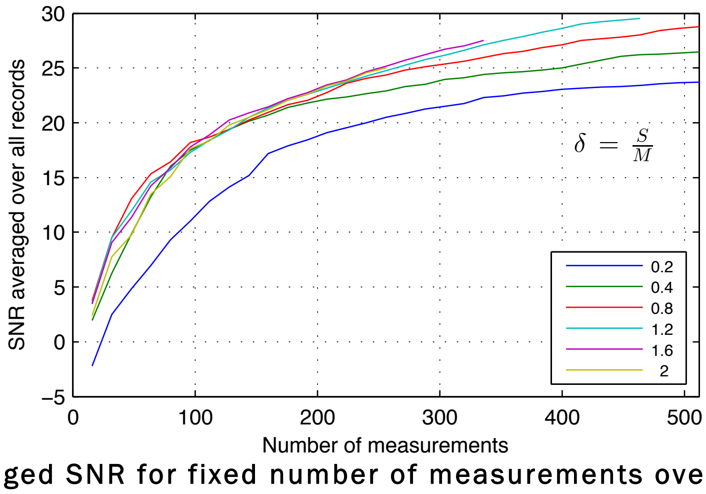

Figure 3: Averaged SNR for fixed number of measurements over different selections of the cardinality of the support, normalized over $M = O(S\log(N/S))$

Based on the results shown on the Figure 3, the best number of for each number of M is extracted. Figure 4 shows the result for best value of S for different number of measurements. The results shows that for small number of measurements, strong similarity exist between the support set of leads, so that a small cardinality in the support is required to reach the best results. Instead, for higher values of the M , similarity between the support set of the leads could not add any additional information. In other words, in case of higher number of measurements, since signals are not exact sparse, even coefficients with a small value impact the quality of the input reconstruction.

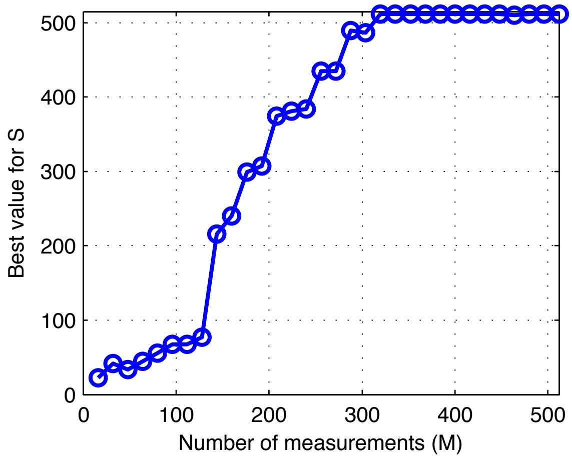

Figure 4: Best value for S corresponding to different values of measurements' number (M).

Figure 5 and Figure 6 shows the comparison between the performance quality of using normal reconstruction for each lead or using the joint compression technique interms of SNR and PRD performance quality respectively. The results shows that expoliting these similaraties between the support set amonge different leads could reach to an important gain, when a small number of measurements is considered.

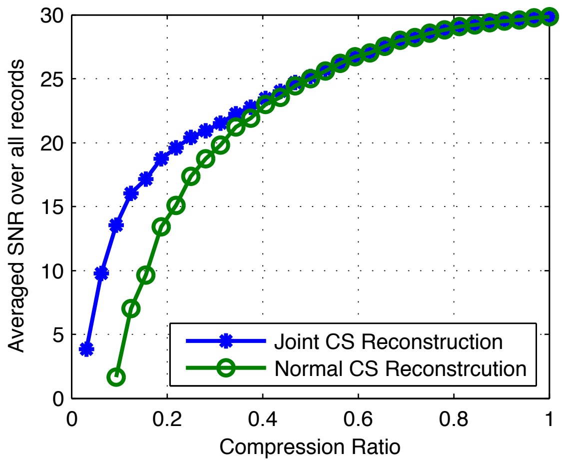

Figure 5: Comparison of Joint CS reconstruction and normal CS

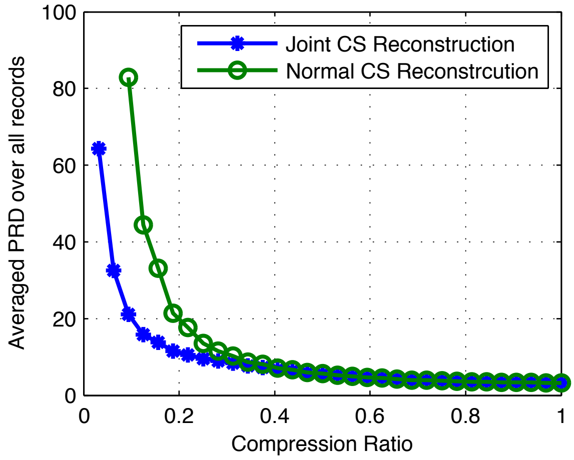

Figure 6:Comparison of Joint CS reconstruction and Normal CS

The results show that by applying the joint reconstruction even a very small number (~20%) of measurements would suffice to reach a good signal reconstruction. Also promising is that even for very small number of measurement (like 5%) where normal CS reconstruction fails, by using the joint CS still it is possible to recover the signals and reach to the performance quality of 10 dB.

Figure 7 also shows the resulting averaged SNR for each individual lead. It can be noted that performance for some of the leads –more specially leads number 4 and 7- are not very, good which affecting the overall SNR. If we compare this result with Normal CS reconstruction the same behavior is seen (Figure 8).

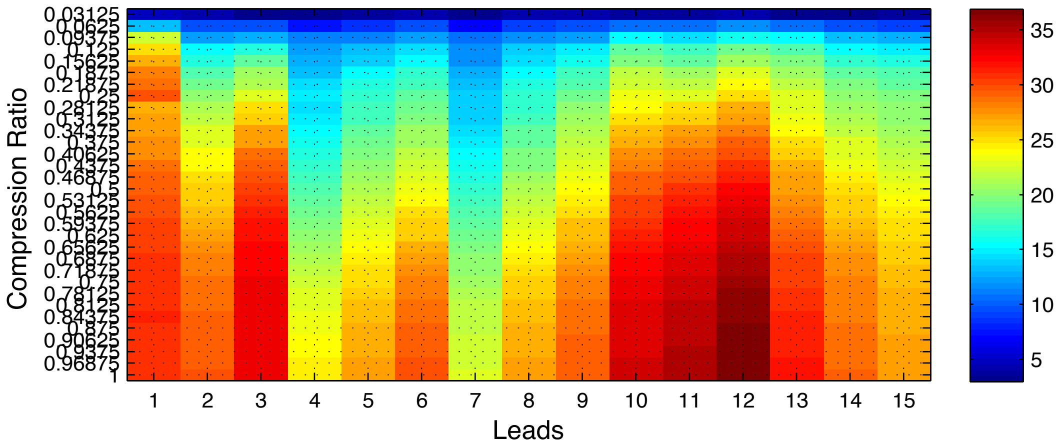

Figure 7:Averaged SNR of Joint reconstruction for each Lead

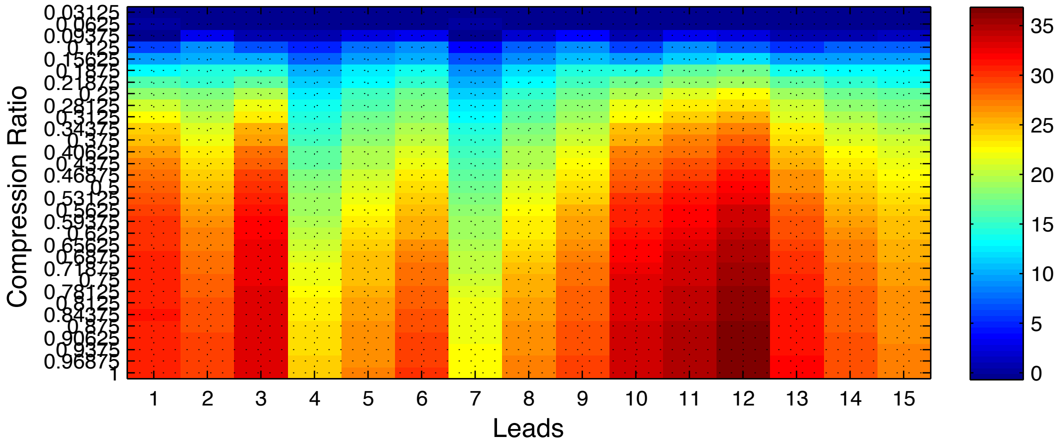

Figure 8: Averaged SNR for normal CS reconstruction for each lead

This is due to the fact that these leads are generally very noisy, explaining the poor behavior of the CS for reconstructing these leads. Figure 9 shows a sample vector of size 512 from these leads. This observation counterintuitively suggests that, even if CS has a poor SNR for some signals, it can be nonetheless beneficial, as it can removing additive noise.

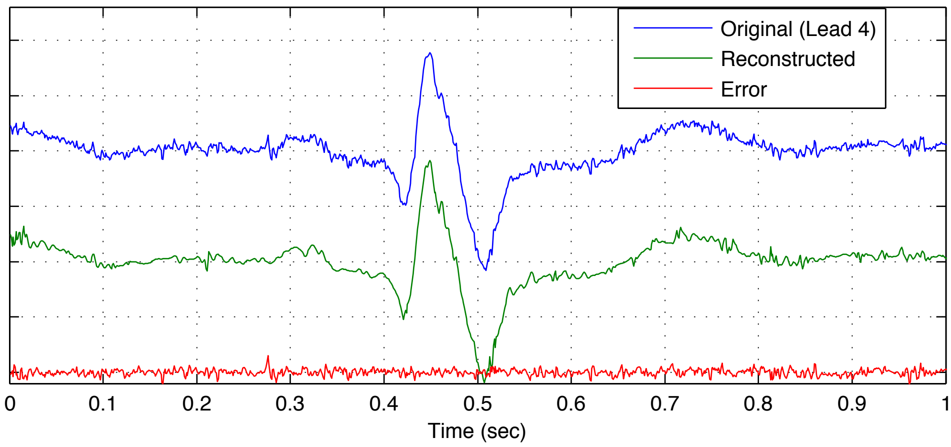

Figure 9: Sample ECG recording for Leads 4 and 7 from PTBDB

Figure 10 and Figure 11 show the box plots comparing the PRD of normal and joint CS reconstruction algorithms. On each box, the central mark is the median, the edges of the box are the 25th and 75th percentiles, and the whiskers extend to the most extreme data points not considered outliers. Results shows that with joint CS reconstruction, performance is increased significantly.

Joint CS reconstruction could reach average signal quality of "Good" reconstruction with CR = 12.5% when for normal CS the required CR is 24%.

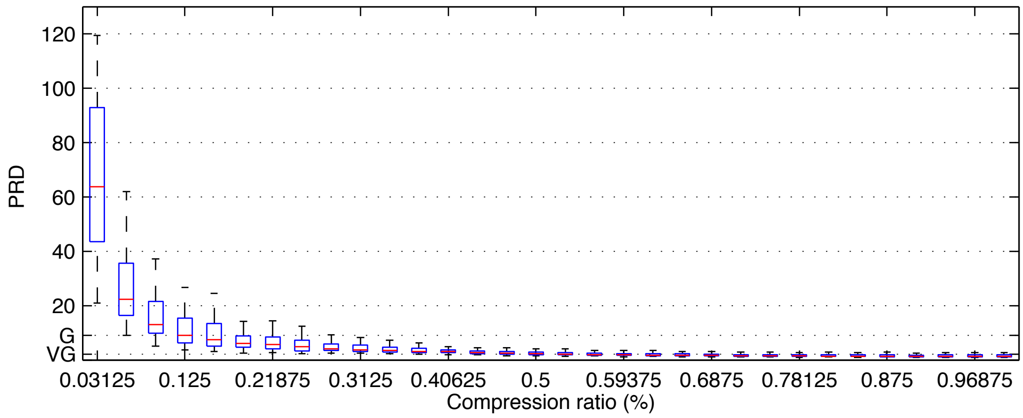

Figure 10: Boxplot for Joint CS reconstruction

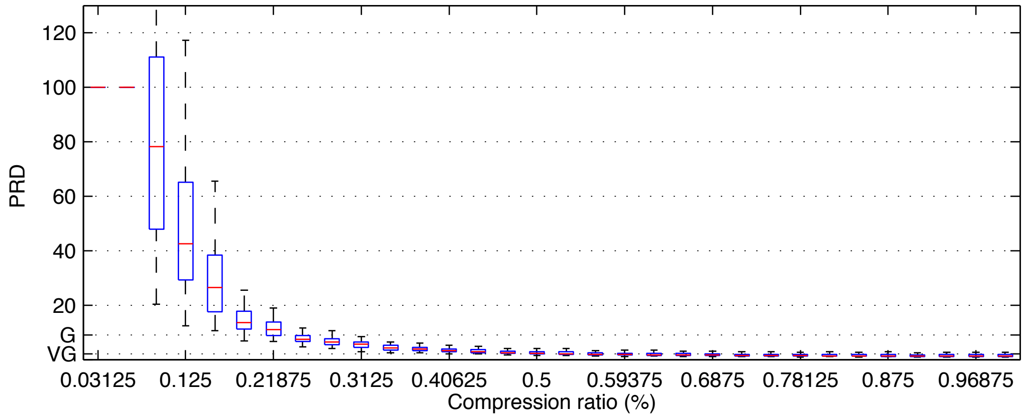

Figure 11: Boxplot For normal CS reconstruction

We have also investigated the reconstruction time for both reconstruction algorithms. Both algorithm are executed on Matlab running of an 15" Macbook with i7 processor chip from intel with OSX 10.6. For each processing black of all records the execution times are recorded and averaged. Figure 12 shows the result of convergence time for both reconstruction algorithms over the different compression ratios. Due to exploiting additional information from other leads the convergence time of the optimization algorithm for joint reconstruction is significantly low compared to the normal CS reconstruction. This is very beneficial and important if the reconstruction algorithm executes on a device with limited amount of resources like an handheld or a mobile device.

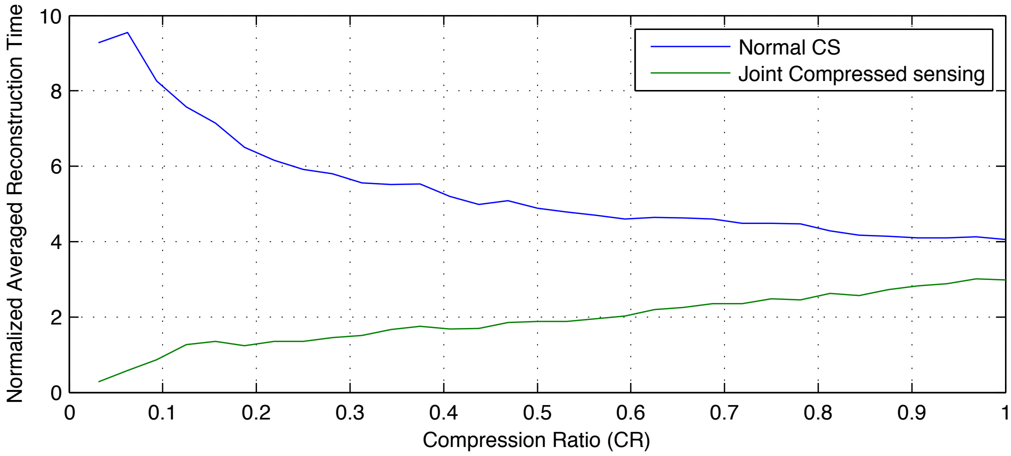

Figure 12: Comparison of averaged normalized reconstruction time

8. Performance of the partners

As lead beneficiary of the deliverable, EPFL investigated the illustrated advanced strategy for CS-based sensing and reconstruction, providing evidence on the similarities in the support of bio-signal acquired concurrently and the benefit of exploiting this information on the quality of reconstructed signals.
All partners fulfilled their tasks in satisfactory time and quality.

9. Conclusions

The deliverable presents advanced an advanced sparsity technique (named joint reconstruction) for compressed sensing and optimized recovery of bio-signals. Joint reconstruction exploits the mutual information present in multi-lead acquisitions to achieve higher compressions of CS-based systems. In the proposed scheme, a lead is sampled at a high compression ratio (CR), will all other leads have a more aggressive, and thus lower, CR. The high-CR signal is used to derive the support (i.e.: index of non-zero coefficients) for all the leads.

We evaluated the performance of joint reconstruction on a database of 15-leads ECG recordings (PTB, available on Physionet). First, we investigated the support of the signals, confirming that strong similarities are present between the leads. Second, we compared the quality of the signal reconstructions when using joint support with respect to a baseline CS implementation. Joint support CS can successfully reconstruct the acquired signals with a lower CR. Moreover, it achieves a better reconstruction quality (lower PRD) for a given CR, especially in the lower range of CR, where a high compression of the signal is performed.

The Full Assembly deems this deliverable to be fulfilled satisfactory.