

Distilling Clinically Interpretable Information from Data Collected on Next-Generation Wearable Sensors

Bryan Haslam, Ankit Gordhandas, Catherine Ricciardi, George Verghese and Thomas Heldt

Abstract—Medical electronic systems are generating ever larger data sets from a variety of sensors and devices. Such systems are also being packaged in wearable designs for easy and broad use. The large volume of data and the constraints of low-power, extended-duration, and wireless monitoring impose the need for on-chip processing to distill clinically relevant information from the raw data. The higher-level information, rather than the raw data, is what needs to be transmitted. We present one example of information processing for continuous, high-sampling-rate data collected from wearable and portable devices. A wearable cardiac and motion monitor designed by colleagues at MIT simultaneously records electrocardiogram (ECG) and 3-axis acceleration to onboard memory, in an ambulatory setting. The acceleration data is used to generate a continuous estimate of physical activity. Additionally, we use a Portapres continuous blood pressure monitor to concurrently record the arterial blood pressure (ABP) waveform. To help reduce noise, which is an increased challenge in ambulatory monitoring, we use both the ECG and ABP waveforms to generate a robust measure of heart rate from noisy data. We also generate an overall signal abnormality index to aid in the interpretation of the results. Two important cardiovascular quantities, namely cardiac output (CO) and total peripheral resistance (TPR), are then derived from this data over a sequence of physical activities. CO and TPR can be estimated (to within a scale factor) from heart rate, pulse pressure and mean arterial blood pressure, which in turn are directly obtained from the ECG and ABP signals. Data was collected on 10 healthy subjects. The derived quantities vary in a manner that is consistent with known physiology. Further work remains to correlate these values with the cardiac health state.

I. INTRODUCTION

Wearable medical devices based on low-cost technology have caught the attention of researchers and companies in the engineering and medical realms. A wide variety of devices is already available, though the field is still growing rapidly, and many challenges remain [1]–[3]. Among other applications, wearable sensors have been developed for heart rate, blood pressure, body position, and glucose monitoring [4], [5]. Engineering challenges that arise universally in the area of wearable medical technology relate to the size and comfort of the device, as well as to power management. Data communication is a central aspect in the design of a wearable

This work was supported in part by a grant from Texas Instruments. B. Haslam is funded by the Department of Defense through a National Defense Science and Engineering Graduate Fellowship.

C. Ricciardi is with the Clinical Research Center, Massachusetts Institute of Technology, Cambridge MA, USA. c.ricci@mit.edu

B. Haslam, A. Gordhandas, G. Verghese and T. Heldt are with the Computational Physiology and Clinical Inference Group, Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge MA, USA. bhaslam@mit.edu, ankitg@mit.edu, verghese@mit.edu, thomas@mit.edu

device, and wireless data transmission consumes substantial power.

A significant challenge also relates to turning the collected raw data into appropriately aggregated clinical information. Since clinicians already face the problem of data overload, wearable devices are unlikely to aid in the clinical decision-making process unless their raw data are converted into clinically interpretable and potentially actionable information. It might therefore be desirable not to stream the collected raw data continuously, but rather to perform much of the data processing on the device itself, and to communicate only clinically important changes in the state of the subject, measured at some aggregate level. The corresponding algorithms will need to be developed to run in real-time on processors that might have limited memory or energy resources. An illustration of this in the setting of electrocardiogram (ECG) waveforms is described in [6].

Since wearable devices are primarily designed for use outside of the relatively well controlled clinical environment, we expect noise and artifacts to be significantly increased in such ambulatory settings. On-chip algorithms are therefore required to extract stretches of data with sufficient signal quality before processing the raw data.

In this paper, we present a method for robust heart rate estimation on the basis of signal-quality assessment, and outline a model-based approach to integrating and interpreting physiological data from different wearable devices. A preliminary version of the results here was presented in [7].

A. Outline

The pumping heart circulates blood through all parts of the body to support metabolism. The average rate of blood flow from the heart is the cardiac output (CO). Various neural, hormonal and metabolic feedback control mechanisms act to maintain arterial blood pressure (ABP) sufficiently high to allow proper perfusion of body tissues, and to keep CO matched to metabolic needs.

Arteriolar resistance is regulated by local metabolic needs, allowing the most metabolically active tissues to receive the greatest blood flow. The total peripheral resistance (TPR) is the net resistance to flow seen by the heart, and is the ratio of mean ABP to CO (in close analogy to electrical resistance, which is the ratio of potential difference to current).

Conventional clinical methods for measuring absolute CO, such as thermodilution or ultrasound, are either invasive or restrictive in methodology, and require expert operators with specialized equipment. Since reliable absolute measurements cannot generally be obtained outside a clinical setting, little

has been done to demonstrate how CO and TPR are affected by normal physical activity.

Using ambulatory sensors, we estimate uncalibrated (i.e., relative, not absolute) versions of both quantities over a series of activities, in addition to obtaining an estimate of the concurrent physical activity of the subject. From these three estimates one can hope to eventually make inferences regarding the relative cardiac health of the subject.

The next section (under Methods) describes our data collection setup and the computations used to estimate CO, TPR and physical activity (PA). We then (under Results) describe representative measured data, as well as the trends in CO, TPR, and PA estimated from this data. Specific discussion of these waveforms and their relationships is contained in the same section, while a broader examination of some issues related to ambulatory monitoring is presented in the last section of the paper (as Discussion).

II. METHODS

A. Devices

The electrical activity associated with the depolarization and repolarization of cardiac muscle on every heart beat is reflected – in a very distinctive and interpretable way – in features of the ECG signal. These features repeat every heart beat, so heart rate is quite easily determined from the ECG signal. More subtle features of the ECG signal can also be used to diagnose a variety of cardiac conditions.

ECG monitoring in the hospital is routine, but finds only occasional use outside the clinic, in the form of Holter monitors. We used an experimental ECG monitor, constructed as a low-noise, low-power device that attaches directly to a set of adhesive ECG electrodes placed on the chest. It can record data for up to two weeks, and perform on-board processing of the data using a Texas Instruments TI MSP 430 processor. There is also an accelerometer on board (Analog Devices ADXL345) for 3-axis accelerometry.

We used a Portapres blood pressure monitoring system to record the ABP waveform continuously. This device is currently the state-of-the-art technology for measuring ABP noninvasively, using a finger cuff. Additionally, we collected clinical ECG data for validation of the experimental device, as well as plethysmography and respiration data.

B. Data Collection

The data collection protocol of our experiments was approved by MIT’s Committee on the Use of Humans as Experimental Subjects. We collected data on 10 subjects. Each subject wore the sensors described in the preceding subsection, and was asked to perform specific tasks mimicking routine ambulatory activities. These included 5 minutes lying down in the supine position, 5 minutes sitting, 5 minutes standing, 1 minute jumping, 1 minute of arm movement, a Valsalva (straining) maneuver, walking and running for 5 minutes each on a treadmill, and stepping for 1 minute. The whole routine was performed twice in succession, with a rest in between, for a total of approximately 1.5 hours per subject.

C. Noise

One significant challenge in working with data collected in ambulatory settings is the level of noise present. We present two approaches to mitigate the effects of noise on the measured data streams.

The first approach leverages the inherent redundancy of simultaneously collected physiological data. More than one of the signals collected contains information about the heart rate of the subject, for example. The ECG is traditionally used to extract heart rate because of the sharp features of the QRS complex, which are easily recognized in each beat and are not unduly susceptible to noise and artifact. The blood pressure waveform, on the other hand, contains the same information, but has lower-frequency fiduciary markers that are susceptible to noise and artifact. Both signals, however, can be used to extract candidate heart rate signals, and the resultant heart rate signals can be compared to each other. Large and sudden deviations in heart rate are usually an indication of poor quality of the underlying data or poor performance of the beat-detection algorithms.

A challenge in processing ECG and blood pressure is that the beat onsets do not occur at the same time (due to the excitation-contraction coupling of the heart and the finite pulse-propagation velocity of the pressure waveform to the measurement location). To identify artifacts in the heart rate signals derived from these two measurements, we developed an algorithm based on the curve-length transform (CLT) to identify large and sudden deviations. For a signal $y(t)$ sampled at intervals Δt , the CLT is defined by

$$\text{CLT}(w, i) = \sum_{k=i}^{i+w} \sqrt{\Delta t^2 + \Delta y_k^2}, \quad (1)$$

where Δy_k is the increment in $y(t)$ between samples $k - 1$ and k , and w denotes the window length, which is a key parameter [8]. The arc length of the curve is calculated over the same three consecutive beats for the two signals. If the resulting values of the CLT differ by a significant fraction (arbitrarily set to 20%), the signal with the lower CLT is used to calculate heart rate for the current beat. If neither signal exhibits high noise, then the average of the corresponding two estimates is taken as the best estimate.

The CLT can be used in a similar fashion to provide a measure of noise in a single signal. A sharp increase in curve length over a particular (well-chosen) window may indicate noise and artifact in that window, rather than physiological change.

D. Modeling and Estimation

Various models and algorithms have been developed in the literature for estimating CO and TPR from more readily accessible quantities such as heart rate and blood pressure. One simple technique is based on the Windkessel model [9]. In this model, the pulse pressure (PP), which is the difference between systolic (peak) and diastolic (valley) pressures on the ABP waveform, is proportional to the stroke volume (i.e., blood volume ejected by the heart at each beat). Since CO is




Fig. 1. Accelerations on the 3 axes are shown in green, cyan and brown, while the physical activity estimate is shown in black. The activities the subject is performing are labeled on the bottom for given stretches of time, with gaps representing rest and setup times.

equal to the stroke volume times the heart rate, CO is then proportional to the product of pulse pressure and heart rate:

$$CO \approx HR * PP . \quad (2)$$

We take this latter product, pulse pressure times heart rate, as our uncalibrated CO estimate. A quantity proportional to TPR can then be estimated by dividing mean ABP by our CO estimate, and this is our uncalibrated TPR estimate:

$$TPR = \text{meanABP}/CO . \quad (3)$$

With the ABP waveform data from the Portapres, we have both the pulse pressure and mean pressure of each beat, as required to compute the above estimates.

To estimate the demands on the cardiovascular system, we track the physical activity (PA) of the subject, using the data from the accelerometer. We estimate PA as the square root of the sum of the variances of all three axes of acceleration, on a sliding four-second window:

$$PA = \sqrt{\sigma_x^2 + \sigma_y^2 + \sigma_z^2} . \quad (4)$$

One limitation of this estimate is that the accelerometer is on the chest, and therefore actually only measures the activity of the torso.

III. RESULTS

Figure 1 shows the accelerometer data along with the associated PA estimate for a portion of the test on one representative subject, as the subject moves through the indicated activities. The remaining figures here correspond to a second subject. The heart rates derived from the ECG and ABP waveforms for this second subject are plotted in Fig. 2, and the filtered, estimated heart rate is plotted in Fig. 3. The resultant estimated heart rate signal does not show many of the unphysiological variations present in Fig. 2. Note from Fig. 3 that two significant drops in HR to 40 beats/min persist. These occurred when the subject performed a Valsalva maneuver, and in fact are a valuable indication of the healthy state of the subject's autonomic control system. The algorithm retained these physiological events, whereas a threshold-based algorithm might have rejected them.




Fig. 2. Noisy heart rate measurements, with ECG-derived heart rate in blue and ABP-derived heart rate in red.




Fig. 3. Cleaned-up heart rate estimate showing very few large deviations. The drops in heart rate around 1600 s and 4200 s are associated with Valsalva maneuvers.




Fig. 4. The cardiac output (blue), total peripheral resistance (green), and physical activity (red) estimates over the course of one experimental session.




Fig. 5. The abnormality score follows the activity estimate well, as anticipated. It can be seen that the second time running on the treadmill was noisier, and a clinician would be directed to the first time running for analysis.

The estimated values for CO, TPR and PA for this subject are plotted in Fig. 4, which shows approximately 90 minutes of data, including two sets of activities described previously. When the subject begins jumping, noted by the sharp rise in PA around 1300 seconds, the CO rises and the TPR falls quickly. After the subject stops, noted by the sharp fall in PA, the CO quickly falls and the TPR rises more gradually.

At around 1800 seconds the subject starts walking on

the treadmill. There is initially a moderate increase in CO, which settles at a level above the previous baseline. The TPR falls, largely reflecting dilation of arterioles in the muscles under local metabolic control, to increase blood flow there. When the subject begins to run on the treadmill around 2100 seconds, the CO rises much more significantly, reflecting the increased metabolic demand. The CO is fairly noisy during this stretch, probably due to the motion of the sensors during running. Immediately after the subject stops running, the CO drops rapidly while the TPR appears to recover more gradually. The interpretation of these post-running transients is more subtle, as it involves the interaction of several mechanisms [10], [11].

At 2700 seconds the subject begins simulated stair climbing by stepping on and off the back of the treadmill. Though the exercise appears to require a similar level of PA as walking, there is a larger increase in CO and an even more gradual decrease in TPR than noted in walking. Afterwards CO falls quickly and TPR rises close to the initial baseline.

The patterns repeat themselves when the subject undergoes the same activities a second time. It may be that during the second set of activities the cardiovascular system reacts faster. This can be seen around 4300 seconds when the subject starts walking on the treadmill and the CO rises faster and higher than it had during the commencement of walking in the first set of activities.

Figure 5 shows a signal abnormality score derived from the CO estimate by simply computing the CLT over a 20-beat window. The noisier portions of the signal, as flagged by higher values of the abnormality score, correspond to periods of higher physical activity.

IV. DISCUSSION

Extracting meaningful clinical information from wearable sensors is a task likely to grow in importance as wearable health monitors make their way out of the research domain and into clinical practice. In this context, two important and related topics need to be addressed: (1) detecting regions of good signal quality for subsequent analysis, and (2) turning raw data into clinically meaningful information. In this paper, we report initial progress on both challenges, in a particular setting. We have developed an adjudication process by which an algorithm determines a robust estimate of heart rate by analyzing differences in heart rate signals derived from two different data streams. The algorithm incorporates a notion of signal quality (or abnormality) to decide which heart rate values to report. Such signal-quality assessment will be very important in handling data from wearable monitors, as motion artifact and other sources of noise are likely to corrupt the signals significantly. We addressed the second challenge by using a model-based algorithm to estimate (uncalibrated) CO and TPR from the acquired ECG and ABP waveforms. Our estimates of CO and TPR show the variations with PA one would expect on the basis of the underlying physiology.

We previously had described results from one subject [7], but in this paper report quite consistent results across all ten subjects tested. The use of the CLT to flag noisy segments

of data, and its use in constructing a more robust estimate of heart rate, were not discussed in [7].

V. CONCLUSIONS

Wearable sensors have not yet made their way into everyday medicine, but have vast potential. We have presented results showing the possibility of estimating important quantities such as CO, TPR, and PA from data collected with wearable and portable sensors. Though not free of noise, the data shows clear patterns that can be explained by the underlying cardiovascular physiology. By using known physiology, we leverage the data we collect from wearable sensors and produce estimates that can be clinically relevant.

VI. ACKNOWLEDGMENTS

Our human subject data collection was conducted in the Clinical Research Center (CRC) at the Massachusetts Institute of Technology, which is supported by NIH Award No. UL1 RR025758, and in part by MIT.

We would like to thank Eric Winokur, Maggie Delano, and Prof. Charles Sodini of MIT's Microsystem Technologies Laboratory for providing the wearable ECG/motion monitor. We are also grateful to Dr. Dennis Buss of TI and Prof. Anantha Chandrakasan of MIT for their interest in and support of this project.

REFERENCES

- [1] A. Pantelopoulou and N. Bourbakis, "A survey on wearable sensor-based systems for health monitoring and prognosis," *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, vol. 40, no. 1, pp. 1–12, 2010.
- [2] A. Lymberis, "Advanced wearable sensors and systems enabling personal applications," in *Wearable and Autonomous Biomedical Devices and Systems for Smart Environment*, ser. Lecture Notes in Electrical Engineering, A. Lay-Ekuakille and S. C. Mukhopadhyay, Eds. Springer Verlag, 2010, vol. 75, pp. 237–257.
- [3] A. Bonfiglio and D. De Rossi, *Wearable Monitoring Systems*. Springer, 2011.
- [4] T. Yilmaz, R. Foster, and Y. Hao, "Detecting vital signs with wearable wireless sensors," *Sensors*, vol. 10, no. 12, pp. 10837–10862, 2010.
- [5] P. Bonato, "Wearable sensors and systems," *Engineering in Medicine and Biology Magazine, IEEE*, vol. 29, no. 3, pp. 25–36, 2010.
- [6] A. Gordhandas, T. Heldt, and G. Verghese, "Real time extraction and analysis of key morphological features in the electrocardiogram, for data compression and clinical decision support," in *AAAI 2011 Spring Symposium on Computational Physiology (SS-11-04)*. AAAI, 2011, pp. 15–18.
- [7] B. Haslam, A. Gordhandas, C. Ricciardi, T. Heldt, and G. Verghese, "Relating noninvasive cardiac output and total peripheral resistance estimates to physical activity in an ambulatory setting," in *AAAI 2011 Spring Symposium on Computational Physiology (SS-11-04)*. AAAI, 2011, pp. 27–31.
- [8] W. Zong, M. Saeed, and T. Heldt, "A QT interval detection algorithm based on ECG curve length transform," *Computers in Cardiology*, vol. 30, pp. 259–262, 2006.
- [9] T. Parlikar, T. Heldt, G. Ranade, and G. Verghese, "Model-based estimation of cardiac output and total peripheral resistance," *Computers in Cardiology*, vol. 34, pp. 379–382, 2007.
- [10] L. Costanzo, *Physiology*, 4th ed. Elsevier, 2010.
- [11] V. Froelicher and J. Myers, *Exercise and the Heart*, 5th ed. Elsevier, 2006.