An Ultra-low Power System-on-Chip for Automatic Sleep Staging

Syed Anas Imtiaz, *Member, IEEE*, Zhou Jiang, *Student Member, IEEE*, and
Esther Rodriguez-Villegas, *Senior Member, IEEE*

Abstract

This paper presents an ultra-low power SoC for automatic sleep staging using a single electroencephalography (EEG) channel. The system integrates an analog front-end for EEG data acquisition and a digital processor to extract spectral features from this data and classify them into one of the sleep stages. The digital processor consists of multiple blocks implementing an automatic sleep staging algorithm that uses a set of contextual decision trees controlled by a state machine. The processor is designed to stay in idle mode at most times waking up only when computations are required. In addition, the mathematical operations are implemented in a way such that the number of datapath components needed is very small. The SoC is implemented in AMS 0.18 μm CMOS technology and is powered using a single 1.25V supply. Its power consumption is measured to be 575 μW while its classification accuracy using real EEG data is 98.7%.

Index Terms

Sleep staging, sleep classification algorithm, low-power biomedical system, electroencephalography, EEG.

S. A. Imtiaz, Z. Jiang and E. Rodriguez-Villegas are with the Circuits and Systems Group, Electrical and Electronic Engineering Department, Imperial College London, UK. Email: ({anas.imtiaz,zhou.jiang11,e.rodriguez}@imperial.ac.uk)). Telephone: +44(0)20 7594 6297.

The research leading to these results has received funding from the European Research Council under the European Community's 7th Framework Programme (FP7/2007-2013) / ERC grant agreement no. 239749.

I. INTRODUCTION

Sleep is a natural state of reduced alertness during which the response of the human body to external stimuli decreases. It is considered a necessity of life for humans and animals alike and is essential to their physical and emotional wellbeing. Sleep accounts for approximately one-third of our lifetime. However, it is not only the number of hours that defines a healthy sleep but also its composition and architecture.

Human sleep is broadly classified into two distinct oscillatory phases based on the eye movements during sleep: Rapid Eye Movement (REM) and Non-Rapid Eye Movement (NREM). According to the American Academy of Sleep Medicine (AASM) rules for sleep classification [1], the NREM phase is further divided into N1, N2, and N3 stages based on the depth of sleep. Hence, together with REM (R) and Wake (W), there are a total five well-defined sleep stages that can be identified from the electroencephalogram (EEG), electrooculogram (EOG), and electromyogram (EMG) signals during sleep. Analysis of these stages, their timing, duration, and occurrence rate are extremely helpful in investigating sleep-related complaints.

In this paper, a fully integrated ultra-low power SoC is presented to acquire data from just a single EEG channel and perform automatic sleep staging. This integrated circuit can be used to create a wearable EEG-based sleep monitoring system. The following section presents an overview of current sleep monitoring technologies with their limitations and challenges and explains how wearable devices can be used to overcome some of these challenges. Section III discusses the possible system architectures for the design of such wearable devices. Section IV briefly describes the sleep staging algorithm that has been developed for an integrated circuit implementation. In Section V, the design and circuit-level implementation of each constituent block of the algorithm is covered in detail. Finally, the SoC performance results and classification accuracy are discussed in Section VI.

II. SLEEP MONITORING AND POLYSOMNOGRAPHY

It is estimated that one third of the global population suffers from abnormal sleep everyday with more than 10% having a clinically significant sleep problem [2]. The symptoms of sleep disorders may include sleep deprivation, disruptive sleep, excessive sleepiness, and other sleep abnormalities that can be fatal if left untreated [3], [4]. These result in social consequences for individuals with sleep problems and a collective financial impact on the economy due to expensive treatments, reduced productivity, accidents, as well as co-morbidities resulting from delayed diagnosis [5].

Sleep disorders are generally presented to the doctors in the form of complaints about difficulties in falling asleep or experiencing unusual sleepiness during daytime. Initial assessments include patient history, interviews, and questionnaires that may be followed up with further tests and monitoring. Patients are normally asked to maintain diaries to log their sleeping habits for about two weeks. These details, together with other lifestyle information, help further the diagnosis and treatment. Although a sleep diary provides invaluable information about the sleeping habits of patients, it is very subjective and based on their perception of sleep. Further, in many cases, more observations and information about the sleep architecture is needed which requires the use of different clinical tools.

A popularly used inexpensive and non-invasive method for sleep monitoring is actigraphy that involves a wrist-worn device to sense movements and record sleep/wake patterns over a long period of time. However, it lacks the ability to provide sleep staging information that is diagnostically vital in many cases [6]. To obtain this, patients are asked to attend a sleep clinic where their sleep is monitored in a test known as polysomnography (PSG). This is the gold standard of sleep monitoring during which brain activity (EEG), eye movements (EOG) and muscle movements (EMG) are monitored together with oxygen saturation, air flow, respiratory effort and others (if needed). The EEG, EOG and EMG signals are manually scored by sleep

technicians in blocks of 30-second epochs based on the AASM rules of classification. The resulting diagnostic graph depicting sleep stages as a function of time is known as a *hypnogram*.

PSG is not very comfortable for patients due to the large number of electrodes attached to them during their sleep. It often requires multiple nights of monitoring to allow time for them to acclimatise and to ensure sufficient data has been obtained. The subsequent manual analysis is also a tedious and error-prone task that can take up to four hours for scoring [7] with high disagreement rate among scorers [8], [9]. All of these make sleep monitoring a significantly time-consuming and costly process. The financial aspect is now more relevant than ever with rising population and soaring healthcare costs globally. As a result, active research is ongoing in the field of sleep medicine to solve some of these problems. To reduce data analysis time and inter-rater disagreements, intelligent algorithms are being developed for automatic sleep stage scoring. There is also an effort to reduce the number of recording channels and to use wireless sensors to make PSG comfortable for patients.

Traditionally, three EEG channels are required in PSG systems together with EOG and EMG channels. However, Ruehland *et al.* [10] reported no significant differences in sleep scoring reliability when using a single EEG channel. The EOG and EMG channels are also traditionally required as REM stage epochs are identified by observing the chin muscle and eye activity [1]. Our recent work, however, demonstrated reliable detection of REM sleep from a single EEG channel which allows using only this channel to perform full sleep staging [11].

While reducing the number of channels, using wireless sensors and adding automatic analysis are important steps forward, in order for a sleep staging system to be practically useful, it must also be small in size and wearable so that the patients can use it with ease and comfort for long hours. At the same time, the analysis must provide reliable results for it to be clinically useful. A low power sleep staging system-on-chip (SoC) can help to meet all these criteria leading to a truly wearable system. With wireless transmission from a single EEG channel and low

power automatic sleep staging, the overall system would be comfortable to use for patients, save analysis time, improve scoring reliability, and reduce overall costs. It could also record sleep parameters more accurately without the need of sleep logs that need to be filled in by patients.

III. WEARABLE SLEEP STAGING SYSTEM

A truly wearable device for automatic sleep staging would entail a wireless non-invasive sensor that transmits data to a small handheld device. Fig. 1 shows two possible architectures for the design of such a system where the only difference between the two is the placement of the signal processing block in the data pipeline.

The first approach uses signal processing at the sensor end so that only the sleep stage information is transmitted every 30 seconds, resulting in a very small data rate. This scenario is an ideal use-case for ultra-low power transmitters, such as Bluetooth Low Energy (BLE) [12] that are widely available in most smartphones. However, due to strict power constraints of the sensing device, its processing capability is very limited. Hence, a trade-off between acceptable levels of performance and algorithm complexity must be made to meet system specifications. In the second approach, signal processing is performed at the receiver end. Thus all raw data needs to be transmitted. With much more relaxed constraints at the receiver, complex algorithms can be used to process the received data. However, using this approach, the transmission stage consumes more power due to a significantly higher data rate.

For both of these approaches, the receiver can either be a handheld device (smartphone, tablet, etc.) or a desktop computer, depending on the processing needs. While the first approach definitely requires low complexity sleep staging algorithms for the aforementioned reasons, limiting the second approach to cases where the receiver is a smartphone also puts some constraints on algorithm development. This is because, compared to a desktop computer or a server, a smartphone will have lower processing resources available. Further, keeping in mind the growing trend of mobile device usage it is safe to say that, at least from the perspective of

user experience, it makes sense to have a system that uses these devices as a receiver. Hence, regardless of the chosen architecture, there is a definite need for low-complexity sleep staging algorithms that can be run on a sensor node or a smartphone. It is for this reason the sleep staging algorithm in [11] was developed. The rest of this paper focuses on the implementation of this algorithm as a complete system-on-chip (SoC).

IV. AUTOMATIC SLEEP STAGING ALGORITHM

A number of sleep staging algorithms have been published in literature over the past five decades, however very few have been implemented in hardware [13], [14]. Most of the reported algorithms have a good classification accuracy which comes at the cost of computational complexity using a large number of features and classifiers. On hardware, this translates to a higher power consumption. Since wearable systems have very limited power budget and computational resources, it is not feasible to run complex algorithms on such systems. Hence, in our previous work [11], a sleep staging algorithm was specifically developed to be used in wearable systems. This is a contextually aware algorithm that uses a set of spectral features and a combination of small decision trees to determine the next sleep stage based on the current stage.

Fig. 2 shows an overview of the sleep staging algorithm. It is based on the observation that when a certain sleep stage is prevalent, it *normally* stays for a few epochs before transitioning to the next stage. This means that the classifier only needs to determine whether an epoch is of the same sleep stage or not using a *one-versus-all* decision tree. If a stage change is detected, a series of *one-versus-one* decision trees are used to determine the new stage. The decision trees, their order of execution and the features needed are established based on the current sleep stage as well as the likelihood of the next one.

There are two levels of tests in this sleep staging algorithm: core tests and peripheral tests. A core test determines whether an epoch being analyzed is of the same sleep stage as the previous one or not. Since there are a total of five sleep stages, the algorithm consists of five core tests

with only one of them active in a given state. If the active core test determines that the current epoch may be of a different sleep stage then a series of peripheral tests are applied to find the new stage. These are very small one-versus-one decision trees that are executed in a specific order. If one of these tests is passed, the rest are not evaluated and the new sleep stage is assigned. In the event that all peripheral tests fail, the state of the machine remains unchanged and the current epoch is assigned the previous sleep stage. There can be a maximum of four peripheral tests corresponding to each core test since the classifier is looking for one of the other four stages except the current stage.

For example, in Fig. 2, the algorithm starts with a previous epoch scored as Wake (this is also the default power-on state). It first checks whether the current epoch is also to be classified as Wake or as one of the other sleep stages using the `Wake vs Others Core Test`. If its result is Wake, no change is needed and the state machine returns to an idle state. If, however, the result is Others, the peripheral tests are used to determine the new state of the machine. For the Wake state, these peripheral tests are: `Wake vs N1`, `Wake vs N2`, `Wake vs N3` and `Wake vs REM`. These will determine the classification label for the current epoch and subsequently result in the state machine being transitioned into the corresponding state. As a result, the next epoch will be classified by starting at a different core test following a similar pattern of peripheral tests.

V. SLEEP STAGING SoC ARCHITECTURE

Fig. 3 shows the four main blocks of the sleep staging SoC: Analog Front-End, Input Controller, Feature Extractor, and Classifier. The function, architecture, and implementation details of these blocks are explained in the following sections.

A. Analog Front-End

The *analog front-end* consists of a low-power and low-noise neural amplifier connected to a 12-bit successive approximation register (SAR) analogue-to-digital converter (ADC) [15]. It uses a clock frequency of 8192 Hz in order to achieve a sampling rate of 256 Hz. The amplifier utilizes an on-chip ultra-high pseudo-resistor ($> 10^{10}\Omega$) to create a high-pass filter with an ultra-low cut-off frequency ($< 0.5\text{Hz}$), thus rejecting large d.c. offsets at its input in a similar configuration to the one in [16]. The neural amplifier achieves a gain of 36 dB while consuming $13.8\ \mu\text{A}$ of current.

B. Input Controller

The sleep staging algorithm processes data in blocks of 2-second subepochs and therefore requires 512 data samples to be buffered prior to any processing. This is performed in the *input controller* which is simply a register bank with an address counter. Shown in Fig. 4, it reads the ADC output at each end of conversion and stores the value at the address pointed by the counter. When 512 samples are received, the counter reaches its maximum limit and is reset to zero. At the same time, a pulse is generated as an instruction for the next block to read the stored data and start computation. This frees up the register bank to continue receiving samples for the next subepoch. The *input controller* also includes circuitry to synchronize the end of conversion signal with the digital clock.

C. Feature Extractor

The *feature extractor* is the computational heart of the sleep staging SoC where most of the mathematical operations take place. At its core is a Fast Fourier Transform (FFT) processor which transforms the signal into frequency domain from which the spectral features are calculated. These features include the spectral edge frequencies and signal power ratios in different frequency bands. More details about these features can be found in [11].

1) *FFT Processor*: The FFT is an algorithm to speed up the computation of the Discrete Fourier Transform (DFT) by making use of its symmetrical properties. This greatly reduces the number of multiplications required, resulting in an algorithmic complexity of $O(N \log N)$ compared to $O(N^2)$ when using the original DFT. The FFT of an input signal x is computed as shown below, where $X(k)$ is the transformed output at index k and N is the length of the signal.

$$X(k) = \frac{1}{N} \sum_{n=0}^{N-1} x(n) e^{-j2\pi nk/N} = \frac{1}{N} \sum_{n=0}^{N-1} x(n) \omega_N^{nk} \quad (1)$$

The most popular algorithm to compute the FFT is the Cooley-Tukey algorithm [17]. Its simplest variant is Radix-2 Decimation in Time (DIT) that breaks down the entire calculation into a number of 2-point DFTs. An N -point FFT requires $N/2$ 2-point DFT computations at each level, with a total of $n(= \log_2 N)$ levels. The mathematical operation involved in computing this DFT is often referred to as the *butterfly* operation and involves one complex multiplication and two complex additions. Hence, a total of $Nn/2$ such butterfly operations are needed to compute an N -point FFT. However, since all butterfly operations are independent of each other only requiring the previous computations to have been completed, a single hardware unit can be used to perform these calculations for different input pairs in a multi-cycle operation provided that the intermediate results from each stage are stored. Consequently, the total number of cycles required to get the final result will be $Nn/2$ if each butterfly operation is performed in one clock cycle.

Fig. 5 shows a block diagram of the FFT processor in the sleep staging SoC. Data samples from the *input controller* are first read and ordered in a bit reversal pattern such that the new position of each sample is determined by flipping the binary bits of its original index. This also ensures the correct order of output samples at the end of the FFT computation. A set of register

banks is used to hold the real and imaginary parts of the initial and intermediate complex values. The address generator keeps track of how many cycles and levels of computations have been performed. Based on the current cycle and level count, it generates addresses for the register banks to fetch and save data for each butterfly computation. The twiddle generator fetches the twiddle factor value from a lookup table for each computation based on the FFT cycle and level. By making use of the symmetry and periodicity properties to extract the twiddle factors, a total of only 64 complex values need to be stored for a 512-point FFT. The butterfly operation is performed each cycle using the fetched twiddle factors and the data points. Its result is saved in-place in the register bank and the fetch-compute-save operation is repeated in the next cycle. When the last of these calculations is complete, the register banks hold the resultant FFT coefficients. The magnitude for each complex output between 0.5-50 Hz is also calculated when the result at that index is available and stored in a separate register bank.

A counter within the FFT processor keeps track of the number of subepochs being processed. This value is passed on to the modules downstream and is used as the subepoch address. A `busy` status signal is raised high while the FFT computation is taking place. Once complete, this is set to low which indicates to the input controller that the FFT process is free to start the next cycle of computation if valid data is available at its input. At the same time, the `flag_fft` signal is set to indicate that the FFT data is available for feature calculation.

2) *Power calculation:* The power calculation module calculates the relative spectral power in different frequency bands using the magnitude values computed by the FFT processor. It normally stays in idle mode and begins computation only after receiving the `flag_fft` signal from the previous module. The implementation of this module is shown in Fig. 6 and its operation involves summing up a set of input magnitude values in a multi-cycle operation. In the first cycle, a register is loaded with the first input from the `data_fft` input bus and an internal counter is incremented. In the subsequent cycles, other inputs are accumulated to this registered value until

the counter reaches the maximum (equal to the total number of inputs). The value in the register at the end is then the power value for the current subepoch and a flag is raised to signify that it is ready to be read. Several instances of this module are used with different FFT magnitudes as input corresponding to the frequency bands in which the power values are to be calculated.

3) *Spectral edge frequency calculation:* The spectral edge frequency (*SEF*) calculation module is used to calculate the *SEF* values at 50% and 95% (*SEF50* and *SEF95*) in various frequency bands. It works by comparing the sum of power in a frequency range against a given threshold. The threshold itself is obtained by multiplying the total power in a given frequency range by the desired percentage (edge). From the hardware point of view, both *SEF50* and *SEF95* modules are the same with one difference. For *SEF50*, multiplication by 0.5 is achieved by shifting one bit to the right. In case of *SEF95*, multiplication by 0.95 is performed using an unsigned fixed-point multiplier with one input set to this constant value.

As shown in Fig. 7, at the start of each *SEF* calculation the power in a given frequency range is obtained using the power calculation module (described in the previous section). This is multiplied by either 0.5 or 0.95 to establish the threshold. At the same time, the maximum frequency value in the given range is also registered. A multi-cycle operation then follows in which the power at the current highest frequency bin is subtracted from the total power. The maximum frequency value is also reduced by 0.5 (frequency resolution) each cycle and updated in the register. When the subtracted power value falls below the established threshold, the corresponding frequency value in the register is the required spectral edge frequency.

4) *Block level implementation of Feature Extractor:* Fig. 8 shows the complete *feature extractor* block with multiple instances of *SEF50*, *SEF95*, and power calculation modules in different configurations to cover the various frequency ranges for all the features. When all the features for an epoch have been calculated, the `flag_fc` status signal is raised to indicate that they are ready to be used for classification.

D. Classifier

The *classifier* is the final stage within the algorithm pipeline which assigns a valid sleep stage to each 30-second epoch of the input EEG signal. It is an implementation of the finite state machine (FSM) approach described in Section IV. It has different states, as shown in Table I, with each state corresponding to a core or peripheral test. In addition, it also has an `IDLE` state during which it waits for valid features to be presented at its input.

When the SoC is powered on, the classifier is in the `IDLE` state. Its next state depends on the score of the last classified epoch and can be one of `CORE_W`, `CORE_N1`, `CORE_N2`, `CORE_N3` or `CORE_R` - each corresponding to one of the core decision trees. The enable signal for the core test corresponding to this new state is asserted low to switch it on from an idle state. The FSM then remains in the same state until all the computations of the core test are complete. If the core test classifies the epoch having the same sleep stage as the previous epoch then the FSM goes back to the initial `IDLE` state since there are no further tests required. If the sleep stage is other than the current one, then the next state of the FSM is one of the peripheral tests. The peripheral tests are also enabled only when needed and their result determines the next state of the FSM until the epoch is assigned a sleep stage.

All core and peripheral tests are similar in design since they take a fixed number of features as inputs, perform one or more comparisons, and produce one of the sleep stages as output. The differences between them include the number of comparisons needed, the input features, and the thresholds corresponding to each feature. This, in turn, dictates how many clock cycles are required by each test to produce the result. Each node in the decision tree performs a comparison of only one feature against its specific threshold. The feature at the node can be of one the following three types:

- *Relative power in a frequency band.* This is the power in a certain frequency band divided by the power in the entire frequency of interest.

- *Ratio of power in different frequency bands.* This is the power in a certain frequency band divided by the power in a different frequency band.
- *Mean spectral edge frequency.* This is obtained by dividing *SEF* by 15 (the number of subepochs in a 30-second epoch).

The three feature types are obtained in a similar fashion by dividing two numbers and then comparing this result against a threshold. This operation can be implemented using a single multiplier/divider and a comparator. At any given time only one of the core or peripheral test is enabled. Within this test, in one clock cycle, only one multiplier and comparator is required across *all* core and peripheral tests since only one such comparison takes place in a single clock cycle. As a result, only one of each component can be shared across the different tests by multiplexing the comparator and multiplier inputs based on which test is currently enabled.

Fig. 9 shows the block diagram of the *classifier* with instances of all core and peripheral tests. Each test is provided with a subset of features that are needed for decision making and an enable signal that activates the test. The tests are connected to a multiplexer that controls the inputs to a single multiplier and a comparator. A state machine controller determines which test needs to be activated at any time, establishes the state of the machine, and also provides the control signal for the multiplexer to allow correct inputs for the shared datapath. The final result of classification is available on the 3-bit output port. It is encoded as shown in Table II and remains valid until the next epoch is available at the classifier input.

VI. SoC TEST RESULTS

A. Measurements

The sleep staging SoC is implemented using the AMS 0.18 μ m process technology with 6 metal layers. The total chip area including the analog front-end 10.3 mm². The digital blocks measure 8.89 mm² in die area with a logic gate count of about 127K. The gain–transfer function

of the amplifier in Fig. 10 shows the gain to be 36 dB in the frequency range of interest. The SoC powered using a single 1.25 V supply and its average power consumption is 575 μ W while operating at a clock frequency of 1.5 kHz. Its performance summary is listed in Table III and a die photograph of the IC is shown in Fig. 11.

The power and area breakdown of the complete SoC are shown in Fig. 12. It can be seen that the digital processing blocks consume about three quarters of the total chip power. They have been synthesized using Cadence RTL Compiler (version 11.21) [18] with several options to reduce their power consumption. These include automatic insertion of clock gates, using slow speed datapath components since the SoC is required to run at a low clock frequency, multi threshold voltage optimization, and insertion of operand isolation. Apart from the tool-aided reduction in power, other methods were also implemented to achieve power savings. Since both core and peripheral tests are seldom activated, the clocks to these tests are gated and enabled when classification using a particular test is needed. Similarly, the clock input to the FFT processor is also gated and switched off when the input controller is collecting data. Further, the algorithm itself was developed with the intent of hardware implementation and allows using small number of computations at the classification stage in order to reach a decision. These design techniques are not exhaustive and other methods such as power gating, multiple supply voltages to lower voltage for certain blocks, and algorithmic improvements can be effectively used to further reduce the power consumption. Additionally, the FFT processor, which consumes about half of the overall SoC power can be improved by following the implementations in [19], [20] and by exploring more efficient architectures [21], [22].

B. System Demonstration

The performance of the the sleep staging SoC was verified using real EEG input while measuring the chip output at intermediate and final stages to ensure operational integrity of

each block. Fig. 13 demonstrates this with a 90-second input EEG signal (equivalent to three 30-second epochs). The algorithm, however, processes data in blocks of 2 seconds, computing the FFT for each block and calculating different features from this result. The acquired data from a subepoch of the first two seconds, followed by a plot its FFT spectrum, as calculated on the chip, is shown in Fig. 13. Using these FFT magnitudes, the different feature values are calculated. A subset of these feature values, together with their Matlab-calculated equivalents, is also shown in Fig. 13 for comparison. Finally, the last plot shows the 90-second output, corresponding to the input EEG. On this plot, the output is invalid during the initial 30 seconds. After this time, the classifier output for the first epoch is available and encoded as 000 (Wake). It becomes invalid briefly before the output for the second epoch appears. This is indicated by a one-bit signal that goes low whenever there is invalid data on the output port.

C. Classification accuracy

The classification accuracy of the sleep staging SoC is determined by comparing its performance against the algorithm which has been developed in MATLAB. The reference algorithm itself used overnight EEG recordings (Fp1-A2 channel) from 20 human subjects in the DREAMS Subjects database [23]. Using the same real EEG signals as input, the overall classification accuracy of the SoC is 98.72% when compared to the reference algorithm. Ideally, this should be 100% but there are a number of factors that makes the hardware different from the reference algorithm. The main difference however is that the reference algorithm uses 64-bit floating-point number representation while the hardware has been designed using 24-bit fixed-point format for representing each number. It is possible to use 64-bit floating-point numbers on the SoC as well but the increase in power consumption and area requirements can be up to four times compared to the fixed-point implementation [24]. Hence, while the use of fixed-point numbers adds truncation and round-off errors leading to misclassified epochs, the reduction in accuracy

is very small compared to the additional resources that would be needed for the floating-point implementation. Further, it must be noted that in most cases a misclassified epoch does not cost the algorithm significantly in terms of the overall accuracy. Once an epoch is incorrectly classified, the state machine controlling the decision trees goes into the wrong state. However, it is expected that the next set of core and peripheral tests will bring it back to the correct state.

VII. DISCUSSION

This paper discusses the design of a complete sleep staging SoC using a single EEG channel. To the best of authors' knowledge, this is also the first implementation of a complete sleep staging algorithm on chip. As a result, it is difficult to make direct comparisons of its performance against other systems that have been designed for different applications. Instead, a performance comparison table of some common circuit blocks that have been used in related state-of-the-art systems is shown in Table IV. However, it should be noted that these systems have been designed for entirely different applications. Regardless, a number of conclusions can be drawn from the work presented in this paper in the context of wearable technologies. The most important conclusion is that having a customized on-chip sleep staging system can result in such low power levels that a one channel EEG system could be running for 370 hours on a typical hearing-aid battery [25]. Considering also that the size of the electronics would also be reduced because of the SoC implementation, and that a typical off the shelf electrode is 10 mm in diameter [26], the overall size of the system can be extremely small. This illustrates how having custom implementation of algorithms on chip can significantly optimize the usability related specifications for wearable technologies, and specifically within the context of this work for sleep monitoring technologies.

Although it could be argued that there are limitations when using just one EEG channel - such as a higher risk of completely losing the signal (associated to the electrode being de-attached), or a reduction in accuracy due to not having EMG or EOG electrodes - the advantages of having just

one channel outweighs the downsides, mostly in the context of certain applications. For example, even when a one channel system may not be a complete replacement for full polysomnography (PSG) carried out in a specialist sleep clinic, it is certainly much more accurate and detailed than actigraphy and can serve as an initial at-home screening device. Hence using just one channel, not only reduces the volume of the system as a whole but also increases its usability (by making it easier to attach). In addition, outside the context of diagnosis, having a long lasting, easy to use, wearable sleep staging system can allow for the possibility of large clinical trials for sleep medicine which would not be possible with traditional PSG systems.

VIII. CONCLUSION

The sleep staging SoC presented here uses an analog front-end for sensing data and a digital processor to run an automatic sleep staging algorithm that has been specifically developed to suit low-power and resource-constrained systems. Its implementation uses a small number of datapath components for feature extraction and classification. Further, all blocks within the system remain in an idle mode waking up only when needed. As a result of its on-chip implementation, the power consumption and area requirements are greatly reduced making the final package small and suitable for wearable use. This demonstrates the potential of our work to be used either as part of a stand-alone sleep monitoring system or as part of a broad-based multi-purpose EEG system on chip that is not limited to sleep monitoring but also useful for other neurological conditions such as epilepsy, stroke or Alzheimers.

REFERENCES

- [1] C. Iber, S. Ancoli-Israel, A. Chesson, and S. Quan, *The AASM manual for the scoring of sleep and associated events: rules, terminology and technical specifications*. Westchester, IL: American Academy of Sleep Medicine, 2007.
- [2] M. Partinen and C. Hublin, "Epidemiology of sleep disorders," in *Principles and Practice of Sleep Medicine*, M. H. Kryger, T. Roth, and W. C. Dement, Eds. Elsevier, 2011, ch. 61.

- [3] H. Yaggi, J. Concato, W. Kernan, J. Lichtman, L. Brass, and V. Mohsenin, "Obstructive sleep apnea as a risk factor for stroke and death," *N. Engl. J. Med.*, vol. 353, no. 19, pp. 2034–2041, 2005.
- [4] J. Horne and L. Reyner, "Sleep related vehicle accidents," *BMJ*, vol. 310, no. 565, 1995.
- [5] G. Stores, "Clinical diagnosis and misdiagnosis of sleep disorders," *JNNP*, vol. 78, no. 12, pp. 1293–1297, 2007.
- [6] J. L. Martin and A. D. Hakim, "Wrist actigraphy," *Chest*, vol. 139, no. 6, pp. 1514–1527, 2011.
- [7] M. Ronzhina, O. Janousek, J. Kolarova, M. Novakova, P. Honzik, and I. Provaznik, "Sleep scoring using artificial neural networks," *Sleep Med. Rev.*, vol. 16, no. 3, pp. 251–63, 2012.
- [8] H. Danker-Hopfe, D. Kunz, G. Gruber, G. Klösch, J. Lorenzo, S. Himanen, T. Kemp, B. amd Penzel, J. Röschke, H. Dorn, A. Schlögl, E. Trenker, and G. Dorffner, "Interrater reliability between scorers from eight european sleep laboratories in subjects with different sleep disorders," *J. Sleep Res.*, vol. 13, no. 1, pp. 63–9, 2004.
- [9] H. Danker-Hopfe, P. Anderer, J. Zeitlhofer, M. Boeck, H. Dorn, G. Gruber, E. Heller, E. Loretz, D. Moser, S. Parapatics, B. Saletu, A. Schmidt, and G. Dorffner, "Interrater reliability for sleep scoring according to the Rechtschaffen & Kales and the new AASM standard," *J. Sleep Res.*, vol. 18, no. 1, pp. 74–84, 2009.
- [10] W. R. Ruehland, F. J. O'Donoghue, R. J. Pierce, A. T. Thornton, P. Singh, J. M. Copland, B. Stevens, and P. D. Rochford, "The 2007 AASM recommendations for EEG electrode placement in polysomnography: impact on sleep and cortical arousal scoring," *Sleep*, vol. 34, no. 1, pp. 73–81, 2011.
- [11] S. Imtiaz and E. Rodriguez-Villegas, "Automatic sleep staging using state machine-controlled decision trees," in *IEEE EMBC*, Milan, August 2015.
- [12] Bluetooth Technology. (2015) Bluetooth Low Energy. [Online]. Available: <http://www.bluetooth.com/Pages/Bluetooth-Smart.aspx>
- [13] J. R. Smith, M. J. Cronin, and I. Karacan, "A multichannel hybrid system for rapid eye movement detection (REM detection)," *Comput. Biomed. Res.*, vol. 4, no. 3, pp. 275–290, 1971.
- [14] J. C. Principe and J. R. Smith, "SAMICOS—a sleep analyzing microcomputer system," *IEEE Trans. Biomed. Eng.*, vol. 33, no. 10, pp. 935–941, 1986.
- [15] P. E. Allen and D. R. Holdberg, *CMOS Analog Circuit Design*. Oxford University Press, 2002.
- [16] R. R. Harrison and C. Charles, "A low-power low-noise CMOS amplifier for neural recording applications," *IEEE J. Solid-State Circuits*, vol. 38, no. 6, pp. 958–965, 2003.
- [17] J. W. Cooley and J. W. Tukey, "An algorithm for the machine calculation of complex Fourier series," *Math. Comp.*, vol. 19, no. 1, pp. 297–301, 1965.
- [18] Cadence. (2015) Encounter RTL Compiler. [Online]. Available: http://www.cadence.com/products/ld/rtl_compiler/pages/default.aspx
- [19] S. R. Sridhara, M. DiRenzo, S. Lingam, S.-J. Lee, R. Blazquez, J. Maxey, S. Ghanem, Y.-H. Lee, R. Abdallah, P. Singh, and M. Goel, "Microwatt embedded processor platform for medical system-on-chip applications," *IEEE J. Solid-State*

- Circuits*, vol. 46, no. 4, pp. 721–730, 2011.
- [20] J. Kwong and A. P. Chandrakasan, “An energy-efficient biomedical signal processing platform,” *IEEE J. Solid-State Circuits*, vol. 46, no. 7, pp. 1742–1753, 2011.
 - [21] Z. Qian and M. Margala, “Low-power split-radix FFT processors using radix-2 butterfly units,” *IEEE Trans. VLSI Syst.*, vol. 24, no. 9, pp. 3008–3012, 2016.
 - [22] M. Garrido, “A new representation of FFT algorithms using triangular matrices,” *IEEE Trans. Circuits Syst. I*, vol. 63, no. 10, pp. 1737–1745, 2016.
 - [23] University of MONS - TCTS Laboratory. (2015) The DREAMS Subjects Database. [Online]. Available: <http://www.tcts.fpms.ac.be/~devuyt/Databases/DatabaseSubjects/>.
 - [24] A. Rushton, *VHDL for Logic Synthesis*, 3rd ed. Chichester, England: John Wiley & Sons, 2011, ch. 15, pp. 337–367.
 - [25] VARTA Microbattery GmbH. (2016) Zinc air wireless approved mercury-free p312. [Online]. Available: <http://www.powerone-batteries.com/en/products/zinc-air-mercury-free-batteries/p312/>
 - [26] Unimed Electrode Supplies. (2016) EEG Electrodes. [Online]. Available: <http://www.unimed-electrodes.co.uk/eeeg>
 - [27] N. Verma, A. Shoeb, J. Bohorquez, J. Dawson, J. Gutttag, and A. P. Chandrakasan, “A micro-power EEG acquisition SoC with integrated feature extraction processor for a chronic seizure detection system,” *IEEE J. Solid-State Circuits*, vol. 45, no. 4, pp. 804–816, 2010.
 - [28] J. Yoo, L. Yan, D. El-Damak, M. A. B. Altaf, A. H. Shoeb, and A. P. Chandrakasan, “An 8-channel scalable eeg acquisition soc with patient-specific seizure classification and recording processor,” *IEEE J. Solid-State Circuits*, vol. 48, no. 1, pp. 214–228, 2013.
 - [29] K. H. Lee and N. Verma, “A low-power processor with configurable embedded machine-learning accelerators for high-order and adaptive analysis of medical-sensor signals,” *IEEE J. Solid-State Circuits*, vol. 48, no. 7, pp. 1625–1637, 2013.
 - [30] W. M. Chen, H. Chiueh, T. J. Chen, C. L. Ho, C. Jeng, M. D. Ker, C. Y. Lin, Y. C. Huang, C. W. Chou, T. Y. Fan, M. S. Cheng, Y. L. Hsin, S. F. Liang, Y. L. Wang, F. Z. Shaw, Y. H. Huang, C. H. Yang, and C. Y. Wu, “A fully integrated 8-channel closed-loop neural-prosthetic cmos soc for real-time epileptic seizure control,” *IEEE J. Solid-State Circuits*, vol. 49, no. 1, pp. 232–247, 2014.
 - [31] M. A. B. Altaf, C. Zhang, and J. Yoo, “A 16-channel patient-specific seizure onset and termination detection soc with impedance-adaptive transcranial electrical stimulator,” *IEEE J. Solid-State Circuits*, vol. 50, no. 11, pp. 2728–2740, 2015.

LIST OF FIGURES

Fig. 1	Two system design approaches for a wireless wearable device: (a) with signal processing at the sensor node; (b) with signal processing at the receiver end.	21
Fig. 2	An overview of the automatic sleep staging algorithm to be implemented on chip.	22
Fig. 3	Top-level block diagram of the automatic sleep staging SoC showing the main modules.	23
Fig. 4	Block diagram of the Input Controller.	24
Fig. 5	Block diagram of the FFT Processor.	25
Fig. 6	Block diagram of the Power calculation module.	26
Fig. 7	Block diagram of the Spectral Edge Frequency calculation module.	27
Fig. 8	Top level diagram of the Feature Calculation block.	28
Fig. 9	Top level diagram of the Classifier block.	29
Fig. 10	Measured gain–transfer function of the amplifier.	30
Fig. 11	Picture of the fabricated Sleep Staging SoC showing analog front-end and digital processor.	31
Fig. 12	(a) Power and (b) Area Breakdown of the sleep staging SoC.	32
Fig. 13	Demonstration of the sleep staging SoC using real EEG signal with intermediate and final output.	33

LIST OF TABLES

TABLE I	Different states of the FSM corresponding to core and peripheral tests. . .	34
TABLE II	Encoding of the output from the classifier block.	35
TABLE III	SoC Performance Summary.	36
TABLE IV	Comparison of SoC Performance with other related systems.	37

Fig. 1. Two system design approaches for a wireless wearable device: (a) with signal processing at the sensor node; (b) with signal processing at the receiver end.

Fig. 2. An overview of the automatic sleep staging algorithm to be implemented on chip.

Fig. 3. Top-level block diagram of the automatic sleep staging SoC showing the main modules.

Fig. 4. Block diagram of the Input Controller.

Fig. 5. Block diagram of the FFT Processor.

Fig. 6. Block diagram of the Power calculation module.

Fig. 7. Block diagram of the Spectral Edge Frequency calculation module.

Fig. 8. Top level diagram of the Feature Calculation block.

Fig. 9. Top level diagram of the Classifier block.

Fig. 10. Measured gain-transfer function of the amplifier.

Fig. 11. Picture of the fabricated Sleep Staging SoC showing analog front-end and digital processor.

(a)

(b)

Fig. 12. (a) Power and (b) Area Breakdown of the sleep staging SoC.

Fig. 13. Demonstration of the sleep staging SoC using real EEG signal with intermediate and final output.

TABLE I
DIFFERENT STATES OF THE FSM CORRESPONDING TO CORE AND PERIPHERAL TESTS.

Core Tests	Peripheral Tests
CORE_W	PERI_W_N1
CORE_N1	PERI_W_N2
CORE_N2	PERI_W_N3
CORE_N3	PERI_W_R
CORE_R	PERI_N1_N2
	PERI_N1_N3
	PERI_N1_R
	PERI_N2_N3
	PERI_N2_R
	PERI_N3_R

TABLE II
ENCODING OF THE OUTPUT FROM THE CLASSIFIER BLOCK.

Sleep Stage	Output Encoding
Wake	000
N1	001
N2	010
N3	011
REM	100
Others/Unknown	111

TABLE III
SoC PERFORMANCE SUMMARY.

Technology	AMS 0.18 μm CMOS (C18A6)
Power Supply	1.25 V
ADC Resolution	12 bits
Gain (dB)	36
CMRR (dB)	62
Logic Gate Count	127K
Total Chip Area	10.3 mm ²
Digital Clock Frequency	1.5 kHz
Analog Clock Frequency	8.192 kHz
Sampling Rate	256 Hz
Classification Accuracy	98.72%
Total Power Consumption	575 μW
Analog Front-End	149 μW
Digital Processor	426 μW

TABLE IV
COMPARISON OF SoC PERFORMANCE WITH OTHER RELATED SYSTEMS.

	Verma [27] JSSC 2010	Yoo [28] JSSC 2013	Lee [29] JSSC 2013	Chen [30] JSSC 2014	Altaf [31] JSSC 2015	This Work
Signal	EEG	EEG	EEG	iEEG	EEG	EEG
Application	Seizure	Seizure	Seizure	Seizure	Seizure	Sleep
# Channels	1	8	1	8	16	1
Feature Extraction	o	o	o	o	o	o
On-chip Classification	x	o	o	o	o	o
AFE Noise RTI	1.3 μ V (0-100Hz)	0.9 μ V (0.5-100Hz)	-	5.23 μ V	0.9 μ V (0.5-100Hz)	2.52 μ V (0.3-130Hz)
Input impedance	> 700 M Ω	> 500 M Ω	-	-	> 500 M Ω	> 500 M Ω
Gain (dB)	40	40	-	41-61	40	36
CMRR (dB)	60	-	-	-	97	62
IA NEF	-	5.12	-	1.77	3.29	5.7
Supply (V)	1.0	1.8/1.0	0.55-1.2	1.8	1.8/1.0	1.25
Classifier	x	Linear SVM	SVM	LLS	D2A-LSVM	Decision Trees
Energy	9 μ J/feature	2.03 μ J/class	273 μ J/class	77.9 μ J/class	2.73 μ J/class	0.7 μ J/class
Technology	0.18 μ m	0.18 μ m	0.13 μ m	0.18 μ m	0.18 μ m	0.18 μ m